

Universal Adversarial Examples in Remote Sensing: Methodology and Benchmark

Yonghao Xu, *Member, IEEE*, and Pedram Ghamisi, *Senior Member, IEEE*

Abstract—Deep neural networks have achieved great success in many important remote sensing tasks. Nevertheless, their vulnerability to adversarial examples should not be neglected. In this study, we systematically analyze the universal adversarial examples in remote sensing data for the first time, without any knowledge from the victim model. Specifically, we propose a novel black-box adversarial attack method, namely Mixup-Attack, and its simple variant Mixcut-Attack, for remote sensing data. The key idea of the proposed methods is to find common vulnerabilities among different networks by attacking the features in the shallow layer of a given surrogate model. Despite their simplicity, the proposed methods can generate transferable adversarial examples that deceive most of the state-of-the-art deep neural networks in both scene classification and semantic segmentation tasks with high success rates. We further provide the generated universal adversarial examples in the dataset named UAE-RS, which is the first dataset that provides black-box adversarial samples in the remote sensing field. We hope UAE-RS may serve as a benchmark that helps researchers to design deep neural networks with strong resistance toward adversarial attacks in the remote sensing field. Codes and the UAE-RS dataset are available online (<https://github.com/YonghaoXu/UAE-RS>).

Index Terms—Adversarial attack, adversarial example, remote sensing, scene classification, semantic segmentation.

I. INTRODUCTION

RECENT advances in remote sensing have brought about the explosive growth of Earth observation data collected by numerous satellite or airborne sensors [1]. The massive availability of these data has significantly boosted the development of many important applications in the geoscience and remote sensing field [2], [3]. Some representative applications include scene classification [4], object detection [5], and semantic segmentation [6]. Currently, most of the state-of-the-art methods for these tasks are based on deep neural networks. On the one hand, deep neural networks—especially convolutional neural networks (CNNs)—have achieved great success in the interpretation of remote sensing data [7]. On the other hand, their vulnerability toward adversarial examples should not be neglected.

In [8], Szegedy et al. first discovered that deep neural networks are very fragile to adversarial examples, which can be simply generated by adding subtle adversarial perturbations to the original image. Such perturbations can be produced via

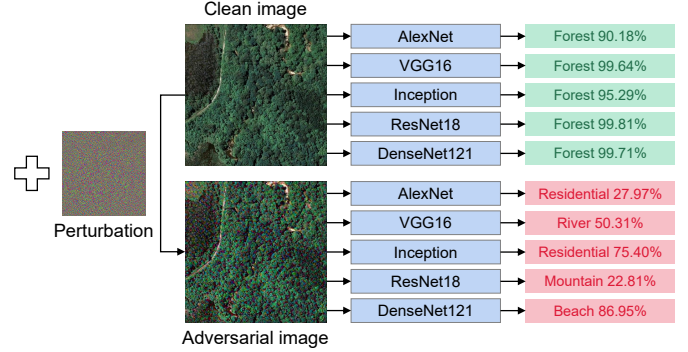


Fig. 1. Illustration of the black-box adversarial attack on the scene classification of very high-resolution remote sensing images using the proposed Mixup-Attack (with ResNet18 as the surrogate model). Without accessing any information from the target model, the generated adversarial image can fool different deep neural networks to make wrong predictions.

specific adversarial attack methods like the box-constrained L-BFGS [8] and fast gradient sign method (FGSM) [9]. With these well-designed algorithms, the generated adversarial examples may look very similar to the original images to the human visual system but can mislead state-of-the-art deep neural networks to make wrong predictions with high confidence [10]. This phenomenon will undoubtedly result in a threat to the security of deep learning-based image recognition systems in real-world applications [11], [12]. For example, Komkov et al. successfully attacked the advanced face identification system using a rectangular paper sticker generated by an algorithm for off-plane transformations [13]. Thus, to improve the reliability of deep learning models, it is necessary to study the characteristics of adversarial examples in depth.

In addition to the aforementioned research that focuses on the computer vision field, there is currently some exploration into adversarial examples in the geoscience and remote sensing field [14]–[16]. Czaja et al. first revealed that adversarial examples also exist in the satellite remote sensing image classification task [17]. Their experiments indicated that adversarial attacks on a small patch inside the remote sensing image could fool deep learning models into making wrong predictions. Chen et al. conducted an empirical study of adversarial examples on scene classification of remote sensing images, where both optical and synthetic-aperture radar (SAR) images are analyzed [18]. Xu et al. further discovered that adversarial examples also exist in the hyperspectral domain [19]. Their experiments revealed that adversarial attacks can successfully change the spectral reflectance characteristics of adversarial hyperspectral samples.

Y. Xu is with the Institute of Advanced Research in Artificial Intelligence (IARAI), 1030 Vienna, Austria (e-mail: yonghao.xu@iarai.ac.at).

P. Ghamisi is with the Institute of Advanced Research in Artificial Intelligence (IARAI), 1030 Vienna, Austria, and also with Helmholtz-Zentrum Dresden-Rossendorf, Helmholtz Institute Freiberg for Resource Technology, Machine Learning Group, 09599 Freiberg, Germany (e-mail: pedram.ghamisi@iarai.ac.at; p.ghamisi@hzdr.de).

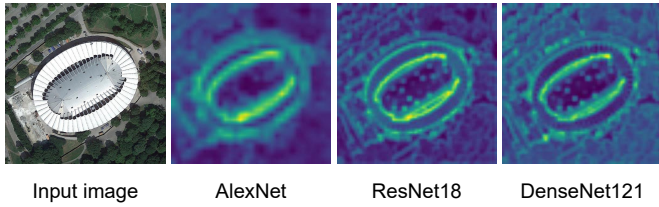


Fig. 2. Illustration of the features in the first pooling layer from AlexNet, ResNet18, and DenseNet121. While different networks have distinct architectures, they may yield similar feature representations in the shallow layers.

So far, existing research on adversarial attacks in the geoscience and remote sensing field mainly focuses on white-box attacks [14], [15], [17], [18], where it is assumed that the attacker has complete knowledge of the victim model, including its architecture and parameter values, so as to generate the corresponding adversarial examples. However, in real-world scenarios, it is usually impossible to obtain detailed information about the deployed network [20]. Under such a circumstance, it is more feasible to conduct a black-box attack, where the adversarial examples are generated without any knowledge of the victim model [10]. Obviously, it is more challenging to implement black-box attacks [21]. One possible solution to achieve this goal is to conduct a white-box adversarial attack with a surrogate model whose complete knowledge is obtainable, and then feed the generated adversarial examples to the unknown target model. Nevertheless, traditional adversarial attack methods like the FGSM are generally designed to fool specific deep neural networks. Thus, the transferability of adversarial examples produced from a particular model to other networks may be limited, especially for those that have architectures that are distinctly different from the model used in the attack. For example, the experimental results in [14] demonstrated that while adversarial examples generated by the AlexNet can significantly cheat AlexNet itself, other deep neural networks like the ResNet or DenseNet possess strong resistance toward this adversarial attack. Therefore, finding the common vulnerability among different networks and generating transferable adversarial examples are critical challenges.

Based on the above analysis, this study aims to conduct black-box attacks for remote sensing data and generate universal adversarial examples that can achieve a high success rate in attacking different deep neural networks without any knowledge about the victim models, as shown in Fig. 1. The initial inspiration of our work comes from an observation that different deep learning models may yield similar feature representations in the shallow layers of the network [22]. Compared to the deep features, which are more abstract, features in the shallow layers generally preserve more detailed spatial information in the image and share similar representations, even in different networks (see Fig. 2 for a visual example). Thus, a natural idea is to implement feature-level adversarial attacks on the shallow layers, which may bring about better transferability to different victim models. To this end, we propose the Mixup-Attack method. Specifically, we first construct the mixup image by the linear combination of images from different

categories. Given a surrogate deep neural network, we extract the shallow features of the mixup image and the target image under attack. Then, we define the mix loss function \mathcal{L}_{mix} by minimizing the KL-divergence between features of the mixup image and the input image. Considering the constraint of \mathcal{L}_{mix} does not take the prediction of the network into consideration, the cross-entropy loss \mathcal{L}_{ce} between the predicted logits on the target image and the true label is also adopted to assist the attack. Thus, the complete objective function \mathcal{L} is the weighted combination of \mathcal{L}_{mix} and \mathcal{L}_{ce} . Finally, the universal adversarial examples can be produced by adding the gradients (adversarial perturbation) of the complete objective function \mathcal{L} to the original clean image. We further propose a variant of Mixup-Attack called the Mixcut-Attack method, in which the mixup image is simply constructed with slices of images from different categories. Despite their simplicity, we find the proposed methods can generate transferable adversarial examples that cheat most of the state-of-the-art deep neural networks without any knowledge about them.

The main contributions of this paper are summarized as follows.

- 1) We systematically analyze the universal adversarial examples in remote sensing data for the first time, without any knowledge about the victim model. Our research reveals the significance of the resistance and robustness of deep learning models when addressing safety-critical remote sensing tasks.
- 2) We propose a novel black-box adversarial attack method, called Mixup-Attack, and its simple variant Mixcut-Attack, for remote sensing data. Extensive experiments on four benchmark remote sensing datasets verify the effectiveness of the proposed methods in both scene classification and semantic segmentation tasks.
- 3) We provide the generated universal adversarial examples in a dataset named UAE-RS, which is the first dataset that provides black-box adversarial samples in the remote sensing field. Experiments demonstrate that samples in UAE-RS can mislead the existing state-of-the-art deep neural networks into making wrong predictions with high success rates, which may serve as a benchmark for researchers developing adversarial defenses.

The rest of this paper is organized as follows. Section II reviews works related to this study. Section III describes the proposed adversarial attack methods in detail. Section IV presents the information on datasets used in this study and the experimental results. Conclusions and other discussions are summarized in Section V.

II. RELATED WORK

A. Adversarial Attacks

1) *Box-Constrained L-BFGS*: The first adversarial attack method was proposed by Szegedy et al., where they generated adversarial perturbations by maximizing the network's prediction error [8].

Formally, let $f : x \in \mathbb{R}^n \rightarrow y \in \mathbb{L}$ be the mapping function of a deep neural network that maps an image with n pixels into a discrete label set. Given an image x and a target label \hat{y} ,

where \hat{y} denotes the wrong label that we expect the network would predict, the adversarial perturbation ρ can be produced by solving the box-constrained optimization problem as below:

$$\min_{\rho} \|\rho\|_2, \text{ subject to : } \begin{cases} f(x + \rho) = \hat{y} \\ x + \rho \in [0, 1]^n. \end{cases} \quad (1)$$

In general, directly solving (1) is a hard problem. Szegedy et al. proposed to approximate the solution of (1) using a box-constrained L-BFGS. More concretely, they performed line-search to find the minimum $c > 0$ for which the minimizer ρ of the following optimization problem satisfies $f(x + \rho) = \hat{y}$, and $\hat{y} \neq y$:

$$\begin{aligned} & \min_{\rho} c\|\rho\|_2 + J(\theta, x + \rho, \hat{y}), \\ & \text{subject to : } x + \rho \in [0, 1]^n, \end{aligned} \quad (2)$$

where θ represents the parameters in the deep neural network, and $J(\cdot)$ denotes the loss function used for training the network (e.g., the cross-entropy loss).

2) *Fast Gradient Sign Method*: In practical applications, the optimization of (2) is still very difficult since it requires layer-wise optimization for parameters in different layers. To make adversarial attacks more efficient, Goodfellow et al. proposed the fast gradient sign method (FGSM) [9]. Given an image x and its true label y , the adversarial example x_{adv} can be calculated as:

$$x_{adv} = \text{clip}(x + \epsilon \text{sign}(\nabla_x J(\theta, x, y))), \quad (3)$$

where $\nabla_x J(\theta, x, y)$ calculates the gradients of the loss function $J(\cdot)$ with respect to the input sample x , $\text{sign}(\cdot)$ denotes the sign function, $\text{clip}(\cdot)$ clips the pixel values in the image, and ϵ is a small scalar value that controls the norm of the perturbation.

Miyato et al. [23] and Kurakin et al. [11] further extend FGSM by applying the ℓ_2 norm and ℓ_∞ norm to the generated perturbation:

$$\ell_2 : x_{adv} = \text{clip}\left(x + \epsilon \frac{\nabla_x J(\theta, x, y)}{\|\nabla_x J(\theta, x, y)\|_2}\right). \quad (4)$$

$$\ell_\infty : x_{adv} = \text{clip}\left(x + \epsilon \frac{\nabla_x J(\theta, x, y)}{\|\nabla_x J(\theta, x, y)\|_\infty}\right). \quad (5)$$

3) *Iterative Fast Gradient Sign Method*: The iterative fast gradient sign method (I-FGSM) was first proposed by Kurakin et al., and is an iterative version of FGSM [11]. At each iteration, the adversarial example can be updated as below:

$$x_{adv}^{t+1} = \text{clip}(x_{adv}^t + \alpha \text{sign}(\nabla_{x_{adv}} J(\theta, x_{adv}^t, y))), \quad (6)$$

where α is the step size. When $t = 0$, x_{adv}^0 is initialized with the original clean image x .

4) *Carlini and Wagner's Attack (C&W)*: Carlini and Wagner proposed to conduct adversarial attacks by encouraging x_{adv} to have a larger probability score for a wrong class than all other classes [24]. This method directly optimizes the distance between the benign examples and the adversarial examples by solving:

$$\arg \min_{x_{adv}} \|x_{adv} - x\|_\infty - \mu J(\theta, x_{adv}, y), \quad (7)$$

where μ is a weighting factor. A more comprehensive review of this method can be found in [25].

B. Black-Box Attacks

Since the discovery of adversarial examples, numerous efforts have been made to design advanced adversarial attack methods [10]. Nevertheless, most of the previous research focuses on the white-box attack, where it is assumed that complete knowledge of the victim model, including its architecture and parameter values, is accessible for the attacker. Although these white-box attack methods can achieve high success rates in the attack, their assumption is usually invalid in practical applications. To make the attack more pragmatic, black-box attack methods are proposed which assume no (or minimal) knowledge about the victim model [10]. Generally, black-box attacks can be categorized into query-based and transfer-based methods.

1) *Query-Based Black-Box Attacks*: Query-based black-box attacks assume that the output of the victim network is accessible during the attack [20], [26]. The main idea is to add subtle perturbations to the input image and observe the output of the victim model. After a series of queries, the real gradients of the victim network can be reconstructed via numerical approximation. For example, Narodytska et al. adopted the greedy local search strategy, where a local neighborhood is used to refine the current image and optimize the objective function related to the network's output in each iteration [27]. To make the query more efficient, Guo et al. proposed to randomly sample a vector from a predefined orthonormal basis and either add or subtract it to the target image [21]. This simple iterative principle can result in higher query efficiency for both targeted and untargeted attacks. Some other representative query strategies include Bayes optimization [28], evolutionary algorithms [29], and meta-learning [30].

2) *Transfer-Based Black-Box Attacks*: The main limitation of query-based black-box attacks lies in the fact that a good approximation for the gradients of a victim model is usually very hard considering the complexity of the deployed model in practice [31]. Thus, it may require a large number of queries to get a high success rate in the attack. By contrast, the main idea of a transfer-based black-box attack is to conduct the white-box attack on a surrogate model and adopt the generated adversarial examples to attack the unknown victim model. Obviously, the success rate of the attack can increase if the surrogate model can share a similar architecture with the victim model. However, since knowledge about the victim model is inaccessible in black-box attacks, directly using traditional attack methods like the FGSM can hardly achieve satisfactory performance. Thus, recent research aims to improve the inherent transferability of the adversarial attack [31]. For example, Liu et al. found that ensemble learning helps to generate more transferable adversarial examples [32]. Zhou et al. and Huang et al. found the attack on the intermediate level is more powerful than the attack on the predicted logits [33], [34]. Chen et al. proposed to conduct adversarial attacks on the attention maps of input images [35].

III. METHODOLOGY

Most of the previous methods focus on attacking the predicted logits of the surrogate model. The essential idea behind

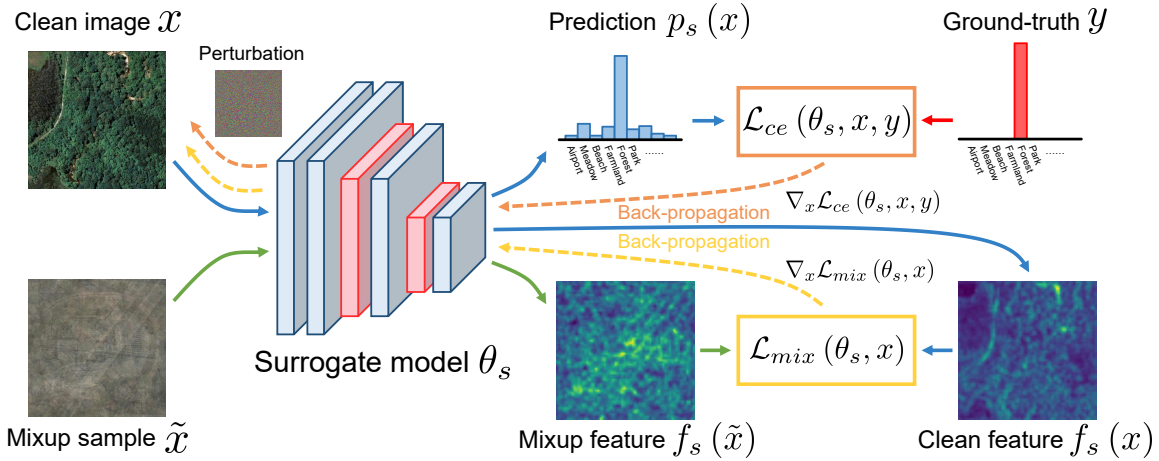


Fig. 3. The proposed Mixup-Attack method for an untargeted black-box adversarial attack on a remote sensing scene classification task. The blue and red blocks represent the convolutional layers and the pooling layers in the surrogate model, respectively.

these methods is to change the pixel values of the input image toward the direction that the surrogate model could yield wrong predicted logits. However, due to the discrepancy between different neural networks, the direction that is harmful to the surrogate model may have limited influence on the victim model [10]. To make the adversarial attack more transferable, we consider attacking the features in the shallow layers of the surrogate model. Since features in the shallow layers generally preserve more detailed spatial information in the image and share similar representations across different networks [22], they may also contain a more common vulnerability. To this end, we first produce a virtual image that belongs to a different category from the input image. Then, instead of directly making the surrogate model yield wrong predicted logits on the input image, we aim to mislead the surrogate model to generate similar feature representations on the input image and the virtual image. Under this circumstance, the generated adversarial examples may also possess stronger transferability to the unknown victim model. Then, a natural question follows: *How can we generate a virtual image whose category can differ from any input image?* To answer this question, this section describes the proposed adversarial attack method in detail.

A. Overview of the Proposed Method

The flowchart of the proposed Mixup-Attack method is shown in Fig. 3. The key idea is to construct a mixup image with the linear combination of training images from different categories and conduct feature-level adversarial attacks on the input image. Specifically, given a surrogate model with parameter θ_s , we extract the shallow features of both the input clean image x and the mixup image \tilde{x} . Then, we define the mix loss function \mathcal{L}_{mix} by minimizing the KL-divergence between features of the mixup image and the input image. Considering the constraint of \mathcal{L}_{mix} does not take the prediction of the network into consideration, the cross-entropy loss \mathcal{L}_{ce} between the predicted logits on the target image and the true label is also adopted to assist the attack. Thus, the complete objective function \mathcal{L} is the weighted combination of \mathcal{L}_{mix} and \mathcal{L}_{ce} .

Finally, the universal adversarial examples can be produced by adding the gradients (adversarial perturbation) of the complete objective function \mathcal{L} to the original clean image. We further integrate the momentum mechanism into the proposed Mixup-Attack to stabilize the update directions in different iterations.

B. Mixup and Mixcut Samples

Our initial inspiration for the proposed Mixup-Attack comes from the work in [36], where the authors proposed a simple and data-agnostic augmentation routine to construct virtual training examples:

$$\begin{aligned}\tilde{x} &= \lambda x_1 + (1 - \lambda) x_2, \\ \tilde{y} &= \lambda y_1 + (1 - \lambda) y_2,\end{aligned}\tag{8}$$

where (x_1, y_1) and (x_2, y_2) are two samples drawn at random from the training set (y_1 and y_2 are one-hot label encoding). $\lambda \in [0, 1]$ is a combination ratio for the mixture.

In fact, (8) extends the training distribution by assuming that linear interpolations of feature vectors should lead to linear interpolations of the associated label vectors. Despite its simplicity, Zhang et al. found training with the augmentation in (8) is very useful in improving the classification performance and adversarial robustness of deep neural networks [36].

While the original design of (8) is to linearly combine two random samples, we extend this mechanism by involving more images from different categories, considering that our goal is to generate the virtual image \tilde{x} whose category can differ from any input image:

$$\tilde{x} = \sum_{i=1}^{n_{mix}} \frac{1}{n_{mix}} x_i,\tag{9}$$

where x_i is a random image from the i th category in the training set ($i = 1, 2, \dots, n_{mix}$), and n_{mix} is the number of categories involved in the Mixup-Attack. Fig. 4 (a) presents the mixup samples with different values of n_{mix} .

Besides the direct linear combination shown in (8), another possible approach is to stitch the slices from different images. Inspired by the work in [37], we further propose the mixcut

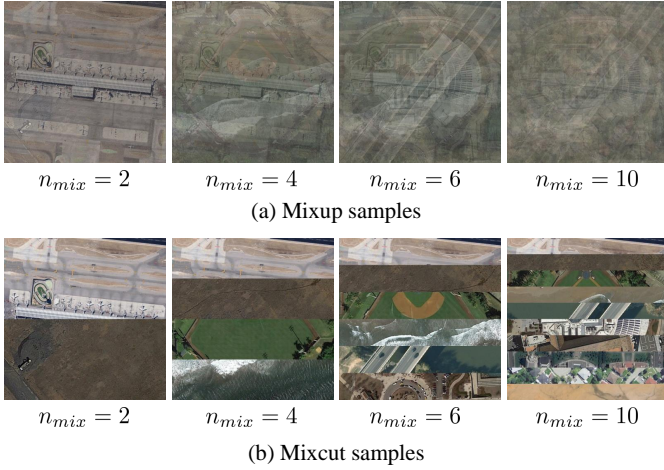


Fig. 4. Illustration of the Mixup and Mixcut samples with different values of n_{mix} .

sampling, which is a simple variant of mixup sampling. Formally, we define the mixcut sampling operation as:

$$\tilde{x} = \sum_{i=1}^{n_{mix}} M_i \odot x_i, \quad (10)$$

where $M_i \in \{0,1\}^{h \times w}$ denotes a binary mask indicating where to drop out and fill in from the i th image x_i ($i = 1, 2, \dots, n_{mix}$), h and w are the height and width of the image, and \odot is element-wise multiplication. For each binary mask M_i , we use the bounding box coordinates $B_i = (r_i, c_i, \Delta r_i, \Delta c_i)$ to indicate the cropping regions, where $r_i = (i-1)h/n_{mix}$, and $c_i = 0$ are the row and column coordinates of the top-left point in the i th bounding box, $\Delta r_i = h/n_{mix}$, and $\Delta c_i = w$ are the height and width of the bounding box. Fig. 4 (b) presents the mixcut samples with different values of n_{mix} .

Note that here we mainly aim to generate a virtual image whose category can differ from any input image; thus, a larger n_{mix} will generally perform better in the attack since there will be more diverse images from different categories involved. We empirically set $n_{mix} = 10$ in our experiments.

C. Feature-Level Attack with Virtual Samples

Recall that our goal is to produce adversarial examples that share similar feature representation with the virtual sample \tilde{x} from (9) or (10). To this end, we propose the mix loss \mathcal{L}_{mix} , which minimizes the distribution discrepancy between the features from the adversarial example and the virtual samples. Formally, given a surrogate model with known parameter θ_s , let $f_s(\cdot)$ denote its mapping functions of feature extraction at the shallow layer. Then, the mix loss $\mathcal{L}_{mix}(\theta_s, x)$ is defined as:

$$\mathcal{L}_{mix}(\theta_s, x) = - \sum_{r=1}^{n_r} \sum_{c=1}^{n_c} \sum_{k=1}^{n_k} f_s(x)^{(r,c,k)} \log \frac{f_s(x)^{(r,c,k)}}{f_s(\tilde{x})^{(r,c,k)}}, \quad (11)$$

where n_r, n_c, n_k are the numbers of rows, columns, and channels in the feature map, respectively. Note that there is

Algorithm 1 Mixup-Attack/Mixcut-Attack

Input:

- A surrogate model θ_s with its feature extraction function f_s and prediction function p_s .
- A clean image x and its ground-truth label y .
- A mixup or mixcut sample generated by (9) or (10).

- 1: $g_0 \leftarrow 0; x_{adv}^0 \leftarrow x$.
- 2: **for** t in $\text{range}(0, T)$ **do**
- 3: Compute the mix loss $\mathcal{L}_{mix}(\theta_s, x_{adv}^t)$, the cross-entropy loss $\mathcal{L}_{ce}(\theta_s, x_{adv}^t, y)$ via (11) and (12).
- 4: Compute the full loss $\mathcal{L}(\theta_s, x_{adv}^t, y)$ via (13).
- 5: Update g_{t+1} by accumulating the velocity vector in the gradient direction:

$$g_{t+1} = g_t + \frac{\nabla_x \mathcal{L}(\theta_s, x_{adv}^t, y)}{\|\nabla_x \mathcal{L}(\theta_s, x_{adv}^t, y)\|_1}.$$

- 6: Update x_{adv}^{t+1} by the normalized gradients of g_{t+1} with ℓ_∞ norm:

$$x_{adv}^{t+1} = \text{clip} \left(x_{adv}^t + \alpha \frac{g_{t+1}}{\|g_{t+1}\|_\infty} \right).$$

- 7: **end for**

- 8: **return** $x_{adv} = x_{adv}^T$.

a negative sign in (11). Thus, maximizing $\mathcal{L}_{mix}(\theta_s, x)$ corresponds to minimizing the KL-divergence between features of the mixup image \tilde{x} and the target image x .

Since the constraint in (11) does not take the prediction of the network into consideration, we further use the cross-entropy loss \mathcal{L}_{ce} to assist the attack. Let $p_s(\cdot)$ denote the mapping functions of the prediction in the surrogate model. Then, the cross-entropy loss $\mathcal{L}_{ce}(\theta_s, x, y)$ is defined as:

$$\mathcal{L}_{ce}(\theta_s, x, y) = - \sum_{k=1}^{n_v} y^{(k)} \log p_s(x)^{(k)}, \quad (12)$$

where n_v denotes the number of categories in the classification task, and y is the true label of the input sample x with the one-hot encoding.

Finally, the complete loss function of the proposed method is defined as:

$$\mathcal{L}(\theta_s, x, y) = \mathcal{L}_{mix}(\theta_s, x) + \beta \mathcal{L}_{ce}(\theta_s, x, y), \quad (13)$$

where β is a weighting parameter for the cross-entropy loss.

With the loss function in (13), the adversarial example x_{adv} can be generated iteratively:

$$x_{adv}^{t+1} = \text{clip} \left(x_{adv}^t + \alpha \frac{\nabla_x \mathcal{L}(\theta_s, x_{adv}^t, y)}{\|\nabla_x \mathcal{L}(\theta_s, x_{adv}^t, y)\|_\infty} \right), \quad (14)$$

where x_{adv}^t is the generated adversarial example at the t th iteration ($t = 0, 1, \dots, T-1$), and α is the step size. When $t = 0$, x_{adv}^0 is initialized with the input image x .

D. Attack with Momentum

Inspired by the work in [40], we further use the momentum strategy to stabilize the update directions at different iterations. Specifically, let g_t denote the momentum term at the t th iteration ($t = 0, 1, \dots, T-1$). When $t = 0$, g_0 is initialized as

TABLE I
SUCCESS RATE (%) OF DIFFERENT UNTARGETED ADVERSARIAL ATTACK METHODS ON THE UCM DATASET.

Surrogate	Method	AlexNet	VGG16	Inception-v3	ResNet18	ResNet101	DenseNet121	DenseNet201	RegNetX-400MF	RegNetX-16GF
AlexNet	FGSM [9]	86.00	13.71	15.24	9.71	7.52	7.14	5.90	11.90	5.24
	I-FGSM [11]	100	34.00	31.62	31.90	22.29	21.33	18.48	35.43	13.90
	C&W [24]	100	33.24	31.14	30.76	22.10	20.86	17.14	33.43	13.52
	TPGD [38]	82.67	27.90	28.95	25.62	16.86	18.57	15.81	26.10	12.10
	Jitter [39]	96.38	13.14	16.10	9.43	7.71	6.67	6.10	10.29	5.14
	Mixup-Attack (ours)	78.86	55.71	56.95	49.24	43.52	31.05	28.10	47.43	40.86
	Mixcut-Attack (ours)	89.05	58.29	58.29	52.10	45.71	35.43	31.43	51.33	42.67
ResNet18	FGSM [9]	17.43	17.81	16.00	67.33	13.05	16.67	15.33	15.90	9.43
	I-FGSM [11]	42.48	48.38	31.71	100	45.24	50.57	46.86	43.52	29.33
	C&W [24]	39.52	46.57	31.24	100	41.33	45.90	45.62	39.81	29.05
	TPGD [38]	37.24	39.05	25.62	75.52	33.33	39.43	33.71	31.05	20.76
	Jitter [39]	21.33	23.43	16.19	92.67	14.38	17.52	18.67	19.24	10.57
	Mixup-Attack (ours)	59.81	71.90	70.86	83.33	49.05	67.81	52.00	50.95	54.86
	Mixcut-Attack (ours)	69.14	80.48	64.29	97.05	71.90	82.86	73.62	72.67	65.14
DenseNet121	FGSM [9]	14.00	14.29	16.48	12.67	10.38	46.48	13.14	11.62	8.00
	I-FGSM [11]	39.62	44.76	30.86	55.33	41.62	97.24	57.05	40.67	35.24
	C&W [24]	36.67	41.24	29.90	53.90	38.57	97.90	55.52	39.33	32.86
	TPGD [38]	35.71	38.00	26.38	44.48	35.71	76.76	44.38	33.71	27.62
	Jitter [39]	18.95	21.62	17.14	18.19	13.71	86.00	21.05	17.52	10.86
	Mixup-Attack (ours)	76.95	74.95	78.76	72.76	70.00	93.14	73.43	70.95	68.76
	Mixcut-Attack (ours)	71.14	75.05	68.10	78.76	73.05	98.95	84.86	74.29	65.14
RegNetX-400MF	FGSM [9]	11.90	10.67	12.57	7.52	7.24	6.86	6.48	60.10	5.43
	I-FGSM [11]	21.14	22.38	18.48	19.33	14.38	17.05	15.81	99.71	17.14
	C&W [24]	20.57	20.76	18.95	19.24	14.10	16.19	14.29	99.62	17.14
	TPGD [38]	19.33	19.05	16.57	15.81	11.90	14.00	12.76	84.48	13.71
	Jitter [39]	12.29	9.43	12.10	7.14	7.14	6.29	6.38	94.48	6.38
	Mixup-Attack (ours)	52.00	49.24	35.24	41.90	35.81	41.14	37.62	85.24	34.76
	Mixcut-Attack (ours)	52.38	50.19	36.92	42.10	36.38	41.05	39.24	86.00	34.86

Note: The leftmost column shows the surrogate models in the adversarial attacks. Best results are highlighted in **bold**.

0. Then, g_{t+1} is updated by accumulating the velocity vector in the gradient direction as:

$$g_{t+1} = g_t + \frac{\nabla_x \mathcal{L}(\theta_s, x_{adv}^t, y)}{\|\nabla_x \mathcal{L}(\theta_s, x_{adv}^t, y)\|_1}. \quad (15)$$

The generated adversarial example x_{adv}^{t+1} at the $(t+1)$ th iteration can thereby be calculated as:

$$x_{adv}^{t+1} = \text{clip} \left(x_{adv}^t + \alpha \frac{g_{t+1}}{\|g_{t+1}\|_\infty} \right). \quad (16)$$

The detailed implementation procedure of the proposed Mixup-Attack and Mixcut-Attack methods are shown in Algorithm 1.

E. Extension to Semantic Segmentation

In this subsection, we generalize the proposed methods to the semantic segmentation of very high-resolution remote sensing images.

We first randomly select 10 images from the training set to generate the mixup/mixcut sample \tilde{x} via (9) or (10). Since semantic segmentation is a pixel-wise classification task, we modify the cross-entropy loss in (12) as:

$$\mathcal{L}'_{ce}(\theta_s, x, y) = - \sum_{r=1}^h \sum_{c=1}^w \sum_{k=1}^{n_v} y^{(r,c,k)} \log \text{up}(p_s(x))^{(r,c,k)}, \quad (17)$$

where h and w are the height and width of the image, and $\text{up}(\cdot)$ denotes the upsampling function with the bilinear interpolation.

Accordingly, the complete loss function of the proposed method is modified as:

$$\mathcal{L}(\theta_s, x, y) = \mathcal{L}_{mix}(\theta_s, x) + \beta \mathcal{L}'_{ce}(\theta_s, x, y). \quad (18)$$

IV. EXPERIMENTS

A. Data Descriptions

1) *Scene Classification*: Two benchmark remote sensing image datasets for scene classification, the UC Merced (UCM)¹ [41] and the AID² [42], are utilized in this study.

UCM consists of 2100 overhead scene images with 21 land-use classes. Each class contains 100 aerial images measuring 256×256 pixels, with a spatial resolution of 0.3 m per pixel in the red-green-blue color space. This dataset is extracted from aerial orthoimagery downloaded from the U.S. Geological Survey (USGS) National Map. The 21 land-use classes are: agricultural, airplane, baseball diamond, beach, buildings, chaparral, dense residential, forest, freeway, golf course, harbor, intersection, medium-density residential, mobile home park, overpass, parking lot, river, runway, sparse residential, storage tanks, and tennis courts.

AID is collected from Google Earth (Google Inc.). It is made up of the following 30 aerial scene types: airport, bare land, baseball field, beach, bridge, center, church, commercial, dense residential, desert, farmland, forest, industrial, meadow, medium residential, mountain, park, parking, playground, pond, port, railway station, resort, river, school, sparse residential, square, stadium, storage tanks, and viaduct. All the images are labeled by specialists in the field of remote sensing image interpretation. The numbers of sample images vary a lot with different aerial scene types, from 220 up to 420. In all, the AID dataset has a number of 10,000 images within 30 classes. The AID dataset has multiple resolutions: the pixel resolution changes from about 8 m to about 0.5 m. The size of each aerial image is fixed at 600×600 pixels.

¹<http://weegee.vision.ucmerced.edu/datasets/landuse.html>

²<https://captain-whu.github.io/AID/>

TABLE II
SUCCESS RATE (%) OF DIFFERENT UNTARGETED ADVERSARIAL ATTACK METHODS ON THE AID DATASET.

Surrogate	Method	AlexNet	VGG16	Inception-v3	ResNet18	ResNet101	DenseNet121	DenseNet201	RegNetX-400MF	RegNetX-16GF
AlexNet	FGSM [9]	88.74	23.32	37.94	19.40	17.74	14.78	11.92	18.62	11.46
	I-FGSM [11]	100	54.82	61.28	54.34	45.22	45.80	36.92	50.96	34.98
	C&W [24]	100	54.04	60.34	53.12	44.80	44.94	35.88	49.60	34.58
	TPGD [38]	81.02	44.26	59.98	42.98	37.30	36.62	27.26	39.30	27.98
	Jitter [39]	97.34	23.64	38.66	20.02	17.60	15.76	11.54	17.36	11.88
	Mixup-Attack (ours)	91.62	82.68	72.28	79.28	69.06	47.18	61.24	74.62	71.40
	Mixcut-Attack (ours)	94.66	84.98	74.56	81.18	71.46	53.14	65.82	77.74	74.48
ResNet18	FGSM [9]	22.52	28.70	38.58	82.16	29.36	31.68	24.22	24.16	22.74
	I-FGSM [11]	47.80	71.18	56.40	99.98	69.02	75.92	71.64	56.06	60.48
	C&W [24]	45.72	65.44	54.90	100	65.04	70.50	65.42	52.90	55.72
	TPGD [38]	40.18	60.50	49.26	75.44	55.84	59.36	54.66	42.98	48.76
	Jitter [39]	24.22	33.52	39.38	96.00	31.48	36.26	27.90	25.10	25.60
	Mixup-Attack (ours)	81.70	75.60	77.00	94.50	92.08	84.86	81.64	76.34	87.04
	Mixcut-Attack (ours)	81.74	86.54	76.52	99.98	90.26	89.84	90.40	80.82	86.66
DenseNet121	FGSM [9]	17.14	20.20	34.92	24.68	21.72	55.04	19.56	18.06	18.24
	I-FGSM [11]	41.42	63.20	51.94	74.58	62.06	99.52	73.94	51.22	56.20
	C&W [24]	40.22	59.42	50.50	69.62	60.54	99.92	70.26	48.86	53.42
	TPGD [38]	38.66	57.52	49.04	66.38	56.18	85.34	64.70	45.38	51.12
	Jitter [39]	21.88	29.16	38.86	37.00	30.48	93.46	30.66	23.36	25.38
	Mixup-Attack (ours)	86.16	88.42	77.60	85.84	90.70	99.12	91.36	85.22	90.24
	Mixcut-Attack (ours)	79.88	87.38	72.78	89.52	87.84	99.98	93.88	83.40	87.08
RegNetX-400MF	FGSM [9]	15.12	19.10	31.76	14.52	14.68	13.28	10.58	60.50	13.46
	I-FGSM [11]	27.94	44.08	43.16	36.08	38.94	34.48	28.60	99.96	36.04
	C&W [24]	27.06	42.62	41.54	34.20	37.26	32.30	26.44	99.82	34.44
	TPGD [38]	24.88	41.12	40.28	30.66	34.68	28.98	22.84	80.22	31.14
	Jitter [39]	16.54	23.24	33.14	15.20	16.50	14.98	11.72	94.26	15.66
	Mixup-Attack (ours)	61.34	81.04	56.28	67.54	74.42	65.84	66.12	92.76	74.24
	Mixcut-Attack (ours)	61.68	81.30	56.62	68.20	75.08	66.58	66.50	92.76	74.80

Note: The leftmost column shows the surrogate models in the adversarial attacks. Best results are highlighted in **bold**.

2) *Semantic Segmentation*: Two benchmark very high-resolution remote sensing image datasets for semantic segmentation, the Vaihingen³ [43] and the Zurich Summer⁴ [44], are utilized in this study.

Vaihingen is a benchmark dataset for semantic segmentation provided by the International Society for Photogrammetry and Remote Sensing (ISPRS); it is a subset of the data used by the German Association of Photogrammetry and Remote Sensing (DGPF) to test digital aerial cameras [43]. There is a total of 33 aerial images with a spatial resolution of 9 cm collected over the city of Vaihingen. The average size of an image is around 2500×1900 pixels, covering an area of about 1.38 km^2 . For each aerial image, three bands are available: the near-infrared, red, and green. Among these images, 16 of them are fully annotated with 6 different land-cover classes: impervious surface, building, low vegetation, tree, car, and clutter/background. We use the same train-test split protocol as specified in the previous work [45] and select five images (image IDs: 11, 15, 28, 30, 34) as the test set. The remaining images are utilized to make up the training set.

Zurich Summer consists of 20 satellite images, taken over the city of Zurich in August 2002 by the Quick-Bird satellite [44]. The spatial resolution is 0.62 m, and the average size of images is around 1000×1000 pixels. The images consist of four channels: the near-infrared, red, green, and blue. Similar to previous work [45], we only utilize the near-infrared, red, and green channels in the experiments and select five images (image IDs: 16, 17, 18, 19, 20) as the test set. The remaining 15 images are utilized to make up the training set. In total, there are 8 urban classes: road, building, tree, grass, bare soil, water, railway, and swimming pool. Uncategorized pixels are

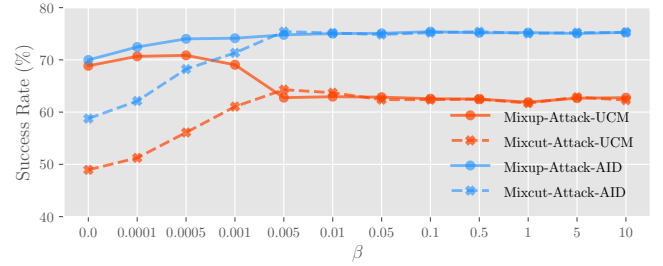


Fig. 5. Success rate (%) of the proposed Mixup-Attack and Mixcut-Attack from ResNet18→Inception-v3 with different values of β on UCM and AID datasets.

TABLE III
PERFORMANCE CONTRIBUTION OF EACH MODULE IN MIXUP-ATTACK AND MIXCUT-ATTACK ON THE SCENE CLASSIFICATION TASK.

Success Rate (%) from ResNet18→Inception-v3				
\mathcal{L}_{ce}	\mathcal{L}_{mix}	Momentum	UCM	AID
<i>Mixup-Attack</i>				
✓			31.71	56.40
✓	✓		64.57	64.54
✓	✓	✓	70.86	77.00
<i>Mixcut-Attack</i>				
✓			31.71	56.40
✓	✓		43.05	61.28
✓	✓	✓	64.29	76.52

Note: Results are reported in success rate (%). Best results are highlighted in **bold**.

labeled as background.

B. Experimental Settings and Implementation Details

We adopt FGSM [9], I-FGSM [11], C&W [24], TPGD [38], and Jitter [39] as the comparison methods, along with the proposed Mixup-Attack and Mixcut-Attack methods, to conduct the untargeted black-box adversarial attack for both

³<http://www2.isprs.org/commissions/comm3/wg4/2d-sem-label-vaihingen.html>

⁴<https://sites.google.com/site/michelevalpiresresearch/data/zurich-dataset>

TABLE IV
SUCCESS RATE (%) OF DIFFERENT UNTARGETED ADVERSARIAL ATTACK METHODS ON THE VAIHINGEN DATASET.

Surrogate	Method	FCN-32s	FCN-8s	DeepLab-v2	U-Net	SegNet	PSPNet	SQNet	LinkNet	FRRNet-A
FCN-8s	FGSM [9]	24.86	22.09	19.17	18.87	20.29	17.88	17.02	17.69	17.31
	I-FGSM [11]	32.39	36.65	24.37	27.59	24.57	25.63	20.49	23.37	25.17
	C&W [24]	42.72	64.28	30.77	38.10	31.62	34.06	24.68	28.60	34.78
	TPGD [38]	26.35	24.41	20.10	20.54	20.80	19.35	17.28	18.56	19.04
	Jitter [39]	25.27	22.55	19.33	19.02	20.39	18.07	17.15	17.86	17.52
	Mixcut-Attack (ours)	37.35	49.65	32.77	47.37	44.31	41.55	35.27	41.14	48.26
	Mixup-Attack (ours)	54.21	65.97	53.92	71.11	71.64	68.85	52.11	65.23	72.05
U-Net	FGSM [9]	24.40	18.41	18.71	23.13	20.17	17.88	16.72	17.80	17.17
	I-FGSM [11]	27.00	21.55	20.44	41.82	22.99	22.83	18.45	20.88	22.80
	C&W [24]	28.06	22.46	20.94	62.85	24.80	23.91	19.09	21.13	26.47
	TPGD [38]	25.40	19.30	19.30	30.90	20.76	18.77	17.05	18.56	19.06
	Jitter [39]	24.55	18.48	18.73	24.09	20.17	18.02	16.83	17.93	17.43
	Mixcut-Attack (ours)	31.68	27.48	23.80	47.74	28.83	29.23	25.25	26.05	29.97
	Mixup-Attack (ours)	46.31	49.41	35.72	58.78	47.17	47.82	47.81	45.81	50.27
PSPNet	FGSM [9]	24.29	18.29	18.60	18.57	19.97	19.78	16.64	17.63	16.81
	I-FGSM [11]	26.47	20.54	19.86	23.04	21.16	32.14	18.07	20.45	20.12
	C&W [24]	27.79	21.25	20.20	24.80	22.02	42.83	18.58	21.11	21.54
	TPGD [38]	25.13	18.80	19.00	19.77	19.95	24.27	16.79	18.33	17.52
	Jitter [39]	24.33	18.25	18.63	18.61	19.91	20.43	16.65	17.72	16.93
	Mixcut-Attack (ours)	28.54	22.90	21.20	26.02	23.31	34.47	21.07	23.18	23.03
	Mixup-Attack (ours)	41.10	43.75	31.75	48.82	45.00	53.41	43.31	44.87	45.86
LinkNet	FGSM [9]	24.79	19.00	19.03	19.54	20.46	18.64	17.59	21.86	17.68
	I-FGSM [11]	28.35	23.50	21.86	26.51	24.15	25.98	21.90	38.23	24.23
	C&W [24]	30.38	25.15	22.85	29.87	27.92	29.53	24.28	55.37	28.14
	TPGD [38]	26.12	20.64	19.94	22.70	21.36	21.55	18.88	29.64	20.53
	Jitter [39]	24.78	18.95	19.02	19.49	20.46	18.66	17.50	22.32	17.69
	Mixcut-Attack (ours)	42.34	40.11	33.92	42.67	41.88	41.75	42.72	52.01	41.14
	Mixup-Attack (ours)	50.74	52.60	40.83	53.56	51.95	47.16	55.62	55.71	49.61

Note: The leftmost column shows the surrogate models in the adversarial attacks. Best results are highlighted in **bold**.

TABLE V
SUCCESS RATE (%) OF DIFFERENT UNTARGETED ADVERSARIAL ATTACK METHODS ON THE ZURICH SUMMER DATASET.

Surrogate	Method	FCN-32s	FCN-8s	DeepLab-v2	U-Net	SegNet	PSPNet	SQNet	LinkNet	FRRNet-A
FCN-8s	FGSM [9]	25.81	13.60	11.75	10.91	13.46	11.71	12.00	10.74	8.88
	I-FGSM [11]	31.54	25.10	19.53	17.57	18.11	16.75	16.02	14.98	15.03
	C&W [24]	43.11	60.56	33.06	28.00	28.78	22.57	22.58	24.27	22.12
	TPGD [38]	27.86	16.33	14.01	13.03	15.75	13.13	12.39	10.91	11.95
	Jitter [39]	27.40	13.61	11.45	11.75	14.52	11.55	12.09	9.38	9.49
	Mixcut-Attack (ours)	47.28	39.83	31.25	24.10	45.15	24.15	27.19	26.56	22.53
	Mixup-Attack (ours)	52.88	46.55	36.66	31.93	52.37	28.29	35.18	33.57	26.32
U-Net	FGSM [9]	25.48	11.05	10.57	13.59	13.69	11.24	11.51	10.59	9.18
	I-FGSM [11]	27.94	15.28	13.95	25.69	17.46	15.04	14.51	13.45	15.17
	C&W [24]	30.84	18.00	16.68	44.91	22.87	16.68	16.63	15.60	17.03
	TPGD [38]	26.98	11.64	11.42	18.84	15.60	11.93	12.24	10.10	12.11
	Jitter [39]	26.65	10.78	10.56	13.45	14.23	11.21	11.77	9.06	9.48
	Mixcut-Attack (ours)	30.38	18.19	18.00	28.99	21.85	19.34	17.68	19.51	19.23
	Mixup-Attack (ours)	37.26	30.98	24.43	52.41	34.42	30.17	31.50	31.70	38.42
PSPNet	FGSM [9]	27.06	11.34	10.95	12.19	14.60	13.44	12.50	9.68	10.06
	I-FGSM [11]	28.68	14.62	13.50	17.81	17.80	26.81	15.94	14.72	16.12
	C&W [24]	29.61	17.00	16.63	20.07	18.90	38.14	18.45	17.99	15.18
	TPGD [38]	26.27	11.50	11.16	11.92	14.88	13.62	11.70	9.99	9.95
	Jitter [39]	26.65	10.70	10.51	11.24	14.06	12.16	11.84	9.08	9.15
	Mixcut-Attack (ours)	29.77	17.83	16.49	22.44	19.61	26.20	17.12	17.92	17.54
	Mixup-Attack (ours)	40.96	34.44	25.51	56.62	33.18	39.13	34.74	38.46	30.76
LinkNet	FGSM [9]	25.48	11.60	12.72	10.93	14.33	13.49	12.77	14.23	10.40
	I-FGSM [11]	30.44	19.16	19.16	19.14	21.46	21.61	20.72	30.64	18.25
	C&W [24]	33.67	23.14	22.14	25.77	26.46	26.32	25.05	48.94	21.62
	TPGD [38]	26.61	13.16	13.34	13.47	15.24	13.90	13.13	15.64	12.43
	Jitter [39]	26.74	11.12	10.89	11.96	14.75	12.84	12.51	12.21	10.91
	Mixcut-Attack (ours)	37.27	25.56	22.94	32.77	39.52	31.95	30.29	37.52	25.32
	Mixup-Attack (ours)	43.36	30.27	27.66	43.45	47.23	38.78	35.92	40.61	29.54

Note: The leftmost column shows the surrogate models in the adversarial attacks. Best results are highlighted in **bold**.

scene classification and semantic segmentation tasks. The perturbation level ϵ and the step size α in all methods are fixed to 1. For iterative methods like I-FGSM, C&W, TPGD, Jitter, and the proposed methods, we fix the number of total iterations T to 5. All methods in this study use the ℓ_∞ norm for the calculation of adversarial perturbation.

The feature extraction function $f_s(\cdot)$ in (11) is implemented with the first pooling layer in each surrogate model. The weighting factor β in (13) and (18) is set as 0.0005 for Mixup-Attack and 0.005 for Mixcut-Attack (see Fig. 5 for the parameter analysis). Scale augmentation [46] is adopted to improve the generalization ability of the model.

TABLE VI
PERFORMANCE CONTRIBUTION OF EACH MODULE IN MIXUP-ATTACK AND MIXCUT-ATTACK ON THE SEMANTIC SEGMENTATION TASK.

Success Rate (%) from FCN8s→SegNet				
\mathcal{L}_{ce}	\mathcal{L}_{mix}	Momentum	Vaihingen	Zurich Summer
<i>Mixcut-Attack</i>				
✓			24.57	18.11
✓	✓		40.92	44.68
✓	✓	✓	44.31	45.15
<i>Mixup-Attack</i>				
✓			24.57	18.11
✓	✓		70.20	52.16
✓	✓	✓	71.64	52.37

Note: Results are reported in success rate (%). Best results are highlighted in **bold**.



Fig. 6. Example images in the UCM dataset (left image in each pair) and the corresponding adversarial examples generated by the proposed Mixcut-Attack method with ResNet18 in the UAE-RS dataset (right image in each pair).

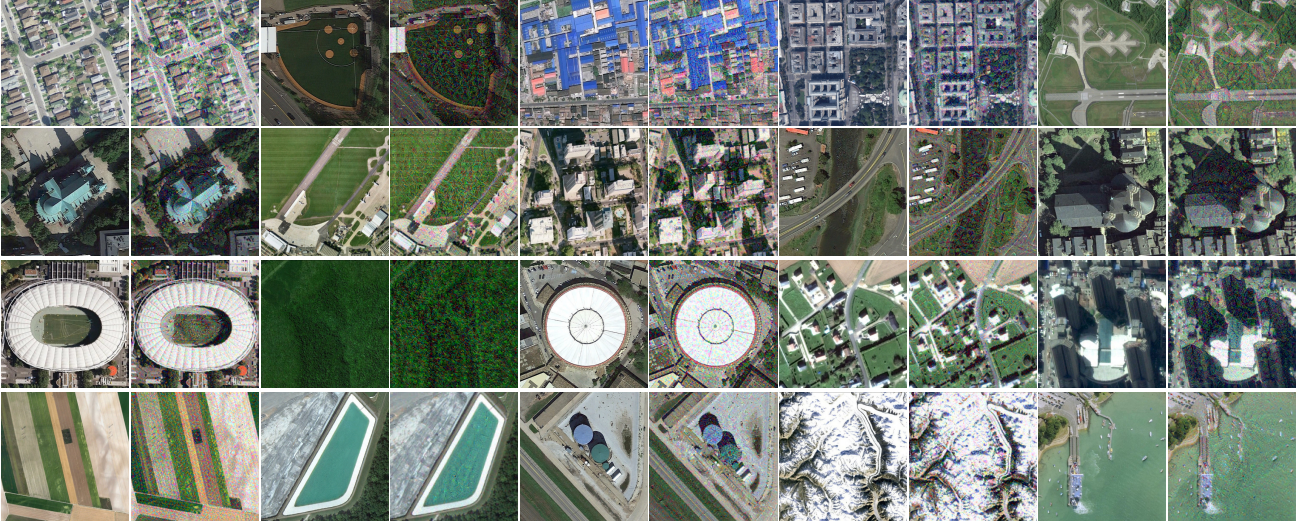


Fig. 7. Example images in the AID dataset (left image in each pair) and the corresponding adversarial examples generated by the proposed Mixcut-Attack method with ResNet18 in the UAE-RS dataset (right image in each pair).

For the scene classification task, we randomly select 50% of the samples for the training set and use the remainder for the test set. We use different surrogate models trained on the training set to generate adversarial examples on the test set with the aforementioned methods, which are then fed to different deep neural networks to evaluate the attack's performance. We adopt the success rate $SR = n_{wrong}/n_{total}$ as the evaluation metric, where n_{wrong} denotes the number of misclassified samples and n_{total} is the number of samples in the test set. A higher success rate indicates that the generated adversarial examples possess stronger transferability to the target model. For the semantic segmentation task, we calculate n_{wrong} by counting the number of all misclassified pixels in the test set, while n_{total} is the sum of all valid pixels in the test set.

The experiments in this study are implemented with the

PyTorch platform [47] using two NVIDIA Tesla V100 (32GB) GPUs.

C. Black-Box Attacks on Scene Classification

Four surrogate models, AlexNet [48], ResNet18 [49], DenseNet121 [50], and RegNetX-400MF [51], are adopted to evaluate the performance of different adversarial attack methods on scene classification.

The detailed quantitative results are presented in Tables I and II. It can be observed that in all black-box attack scenarios where the target model is different from the surrogate model, the proposed Mixup-Attack and Mixcut-Attack can outperform the comparison methods by a large margin in both datasets. Take the results of ResNet18→Inception-v3 on the UCM dataset, for example. While C&W obtains a success rate of about 31%, both Mixup-Attack and Mixcut-Attack can

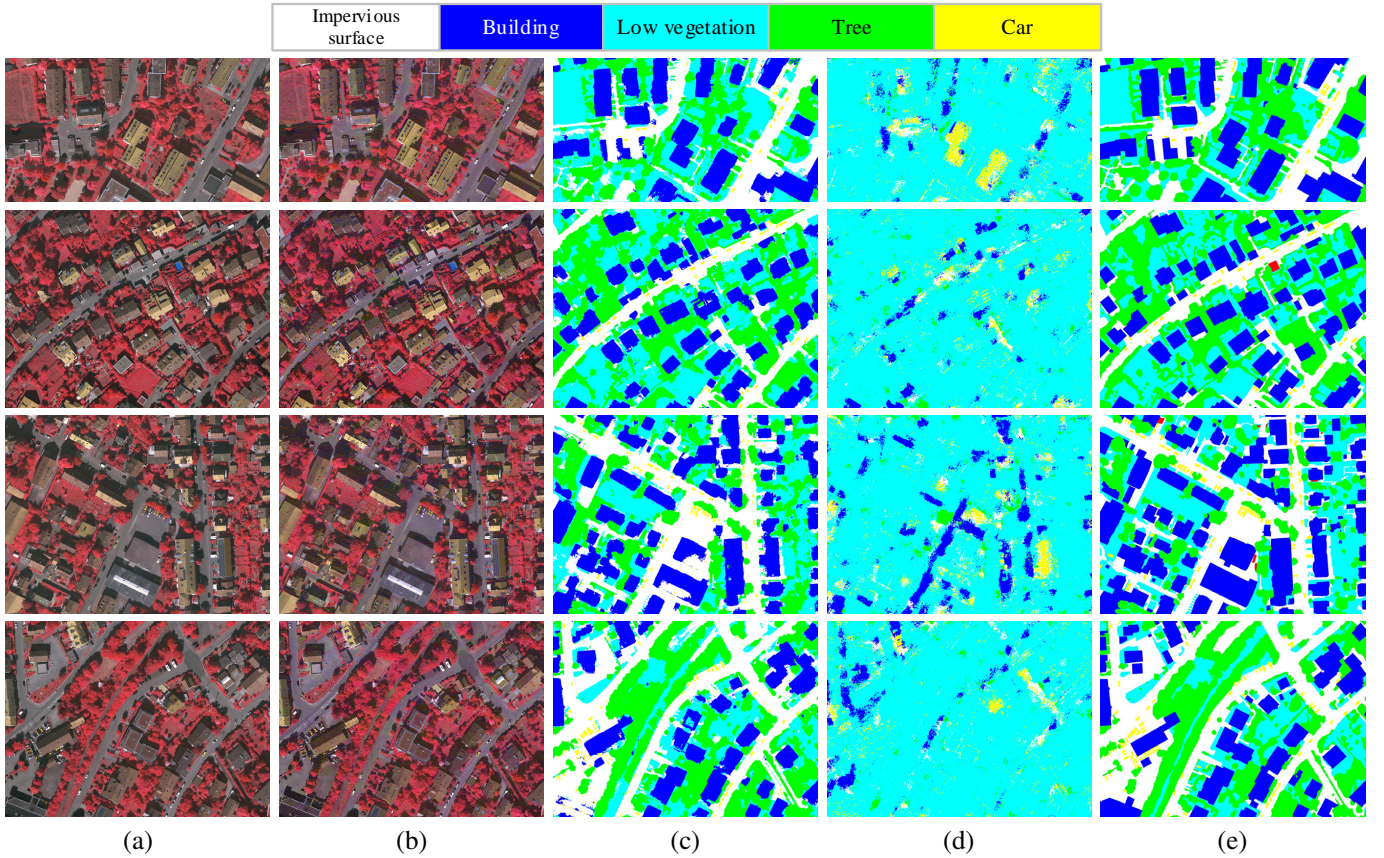


Fig. 8. Qualitative results of the black-box adversarial attacks from FCN-8s \rightarrow SegNet on the Vaihingen dataset using the proposed Mixup-Attack method. (a) The original clean test images in the Vaihingen dataset. (b) The corresponding adversarial examples in the UAE-RS dataset. (c) Segmentation results of SegNet on the clean images. (d) Segmentation results of SegNet on the adversarial images. (e) Ground-truth annotations.

achieve a success rate of over 64%, which is much higher with more than 30 percentage points. Similar phenomena can be observed in the AID dataset. These results indicate that the proposed methods can generate adversarial examples with stronger transferability to different target models.

It can also be observed that the white-box adversarial attack performance of the proposed methods is generally lower than that of I-FGSM or C&W methods. Take the results of ResNet18 \rightarrow ResNet18 in the UCM dataset, for example. While both I-FGSM and C&W methods obtain a success rate of 100%, Mixup-Attack yields a success rate of about 83% in this case. One intuitive reason for this phenomenon is that the mix loss $\mathcal{L}_{mix}(\theta_s, x)$ used in the proposed methods may sacrifice a little bit of white-box attack performance to achieve better transferability, as it does not directly mislead the network to yield wrong predicted logits. Nevertheless, in most cases, the proposed methods can still obtain a success rate over 80% in both datasets, which would be a serious threat in the white-box adversarial attack scenario.

The parameter β in (13) is an important factor in the proposed methods. Fig. 5 shows the success rate of the proposed Mixup-Attack and Mixcut-Attack from ResNet18 \rightarrow Inception-v3 with different values of β on UCM and AID datasets. It can be observed that for Mixup-Attack, a relatively small β would be more beneficial to the performance, especially on the UCM dataset, where the highest success rate is obtained

with $\beta = 0.0005$. For Mixcut-Attack, it can be observed that the $\mathcal{L}_{ce}(\theta_s, x, y)$ term plays an important role in the model, as the success rate increases accordingly as the value of β grows on both datasets. When $\beta > 0.005$, the success rate gradually becomes saturated. Based on the above analysis, we set $\beta = 0.0005$ for Mixup-Attack and $\beta = 0.005$ for Mixcut-Attack in the experiments.

To evaluate how each module in the proposed methods would influence the adversarial attack performance, the quantitative ablation study results are presented in Table III. In both UCM and AID datasets, we find that directly using the cross-entropy loss \mathcal{L}_{ce} alone only leads to a limited success rate, while combining \mathcal{L}_{ce} and \mathcal{L}_{mix} can significantly improve the performance by around 33 and 8 percentage points for Mixup-Attack, and 12 and 5 percentage points for Mixcut-Attack. Finally, with the help of the momentum strategy, the success rate is further increased, achieving state-of-the-art performance.

D. Black-Box Attacks on Semantic Segmentation

Four surrogate models, the FCN-8s [52], U-Net [53], PSP-Net [54], and LinkNet [55], are adopted to evaluate the performance of different adversarial attack methods on semantic segmentation.

The detailed quantitative results are presented in Tables IV and V. Compared to the scene classification models, the

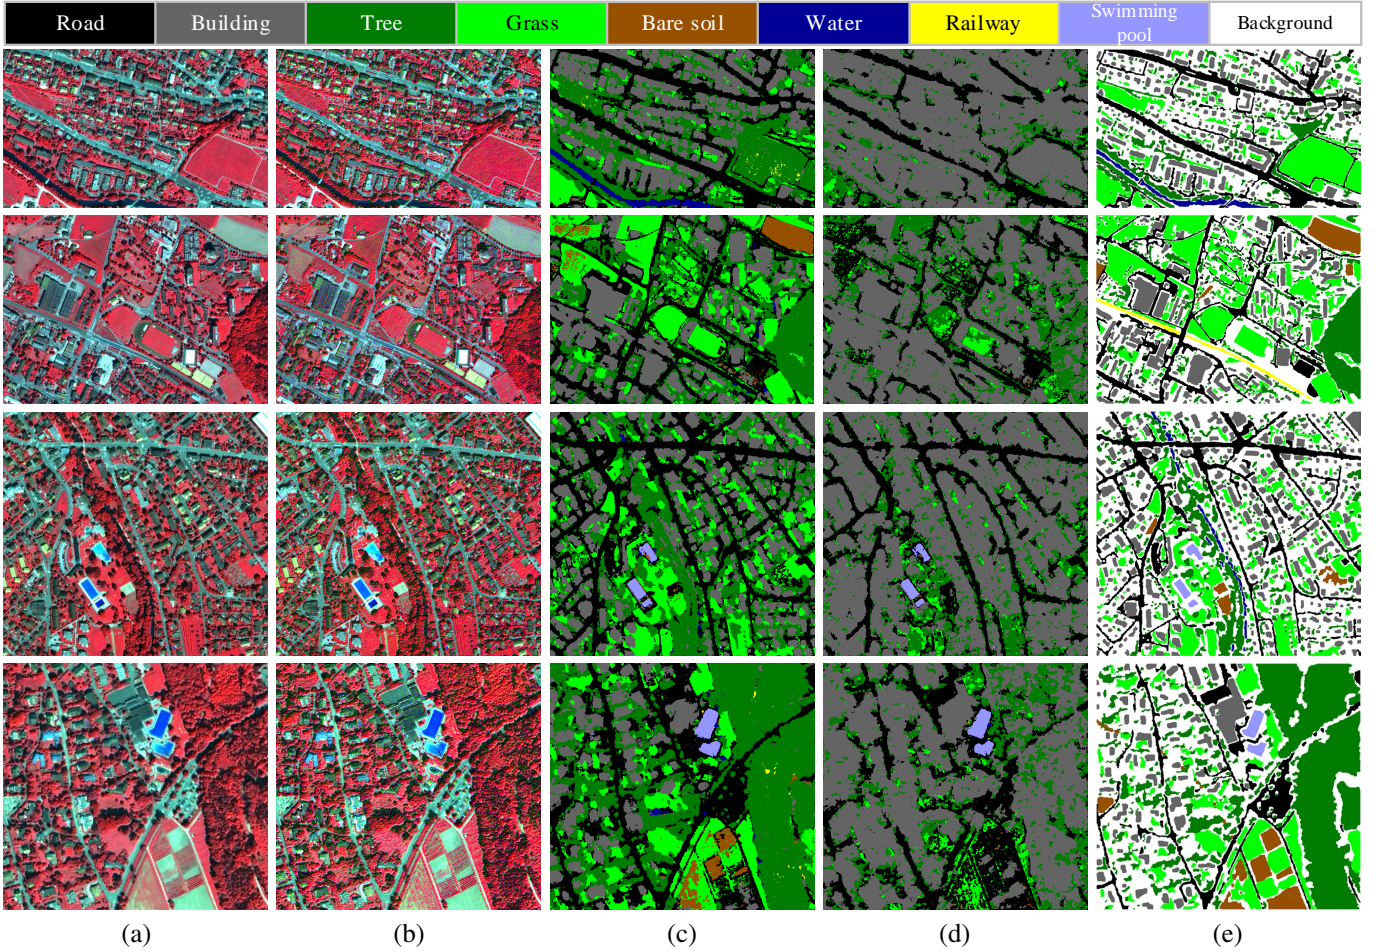


Fig. 9. Qualitative results of the black-box adversarial attacks from FCN-8s \rightarrow SegNet on the Zurich Summer dataset using the proposed Mixup-Attack method. (a) The original clean test images in the Zurich dataset. (b) The corresponding adversarial examples in the UAE-RS dataset. (c) Segmentation results of SegNet on the clean images. (d) Segmentation results of SegNet on the adversarial images. (e) Ground-truth annotations.

semantic segmentation models are generally more difficult to attack, as pixel-wise adversarial perturbations are required for the latter models. Nevertheless, it can be observed from Tables IV and V that in all black-box attack scenarios, the proposed Mixup-Attack can still outperform the comparing attack methods by large margins. Take the results of FCN-8s \rightarrow SegNet on the Vaihingen dataset, for example. While C&W obtains a success rate of about 38%, the proposed Mixup-Attack can achieve a success rate of over 71%, which is much higher with more than 30 percentage points. Similar phenomena can be observed in the Zurich Summer dataset.

It can also be observed that while Mixcut-Attack can achieve competitive performance when attacking scene classification models, its success rate is relatively lower than that of Mixup-Attack for the semantic segmentation task. One possible reason is that mixcut samples only perturb the real spatial distribution of different pixels in the clean image but cannot change the pixel values. Thus, they could be helpful in fooling scene classification models where only one label is predicted for a whole image but would have less influence on semantic segmentation models, as they would make a pixel-wise classification for the whole image.

To evaluate how each module in the proposed methods

influences the adversarial attack performance for the semantic segmentation task, the quantitative ablation study results are presented in Table VI. In both Vaihingen and Zurich Summer datasets, we can find that directly using the cross-entropy loss \mathcal{L}_{ce} alone can hardly achieve a high success rate, while combining both \mathcal{L}_{ce} and \mathcal{L}_{mix} can significantly improve the performance by around 45 and 34 percentage points for Mixup-Attack, and 16 and 26 percentage points for Mixcut-Attack. Besides, the momentum strategy brings about relatively limited success rate improvement which is different from the scene classification scenario in Table III.

E. UAE-RS Dataset

Considering the high success rate of the proposed methods for the black-box adversarial attack, we collect the generated adversarial examples in a dataset named UAE-RS. To the best of our knowledge, UAE-RS is the first adversarial examples dataset in the remote sensing community, which may serve as a benchmark that helps researchers to design deep neural networks with strong resistance toward adversarial attacks.

To build UAE-RS, we use the Mixcut-Attack method to attack ResNet18 with 1050 test samples from the UCM

TABLE VII
QUANTITATIVE SCENE CLASSIFICATION RESULTS OF DIFFERENT DEEP NEURAL NETWORKS ON THE CLEAN AND UAE-RS TEST SETS.

Model	UCM			AID		
	Clean Test Set	UAE-RS Test Set	OA Gap	Clean Test Set	UAE-RS Test Set	OA Gap
AlexNet [48]	90.28	30.86	-59.42	89.74	18.26	-71.48
VGG11 [56]	94.57	26.57	-68.00	91.22	12.62	-78.60
VGG16 [56]	93.04	19.52	-73.52	90.00	13.46	-76.54
VGG19 [56]	92.85	29.62	-63.23	88.30	15.44	-72.86
Inception-v3 [57]	96.28	24.86	-71.42	92.98	23.48	-69.50
ResNet18 [49]	95.90	2.95	-92.95	94.76	0.02	-94.74
ResNet50 [49]	96.76	25.52	-71.24	92.68	6.20	-86.48
ResNet101 [49]	95.80	28.10	-67.70	92.92	9.74	-83.18
ResNeXt50 [58]	97.33	26.76	-70.57	93.50	11.78	-81.72
ResNeXt101 [58]	97.33	33.52	-63.81	95.46	12.60	-82.86
DenseNet121 [50]	97.04	17.14	-79.90	95.50	10.16	-85.34
DenseNet169 [50]	97.42	25.90	-71.52	95.54	9.72	-85.82
DenseNet201 [50]	97.33	26.38	-70.95	96.30	9.60	-86.70
RegNetX-400MF [51]	94.57	27.33	-67.24	94.38	19.18	-75.20
RegNetX-8GF [51]	97.14	40.76	-56.38	96.22	19.24	-76.98
RegNetX-16GF [51]	97.90	34.86	-63.04	95.84	13.34	-82.50

Note: Results are reported in overall accuracy (%).

TABLE VIII
QUANTITATIVE SEMANTIC SEGMENTATION RESULTS OF DIFFERENT DEEP NEURAL NETWORKS ON THE CLEAN AND UAE-RS TEST SETS.

Model	Vaihingen			Zurich Summer		
	Clean Test Set	UAE-RS Test Set	Mean F_1 Gap	Clean Test Set	UAE-RS Test Set	Mean F_1 Gap
FCN-32s [52]	69.48	35.00	-34.48	66.26	32.31	-33.95
FCN-16s [52]	69.70	27.02	-42.68	66.34	34.80	-31.54
FCN-8s [52]	82.22	22.04	-60.18	79.90	40.52	-39.38
DeepLab-v2 [59]	77.04	34.12	-42.92	74.38	45.48	-28.90
DeepLab-v3+ [60]	84.36	14.56	-69.80	82.51	62.55	-19.96
SegNet [61]	78.70	17.84	-60.86	75.59	35.58	-40.01
ICNet [62]	80.89	41.00	-39.89	78.87	59.77	-19.10
ContextNet [63]	81.17	47.80	-33.37	77.89	63.71	-14.18
SQNet [64]	81.85	39.08	-42.77	76.32	55.29	-21.03
PSPNet [54]	83.11	21.43	-61.68	77.55	65.39	-12.16
U-Net [53]	83.61	16.09	-67.52	80.78	56.58	-24.20
LinkNet [55]	82.30	24.36	-57.94	79.98	48.67	-31.31
FRRNet-A [65]	84.17	16.75	-67.42	80.50	58.20	-22.30
FRRNet-B [65]	84.27	28.03	-56.24	79.27	67.31	-11.96

Note: Results are reported in mean F_1 score (%).

dataset and 5000 test samples from the AID dataset for scene classification, and use the Mixup-Attack method to attack FCN8s with 5 test images from the Vaihingen dataset (image IDs: 11, 15, 28, 30, 34) and 5 test images from the Zurich Summer dataset (image IDs: 16, 17, 18, 19, 20) for semantic segmentation.

1) *UAE-RS Samples on Scene Classification*: Some example images in the original UCM and AID datasets and the corresponding adversarial images in the UAE-RS dataset are presented in Figs. 6 and 7. It can be observed that the generated adversarial images may look very similar to the original clean images to a human observer, despite their ability to seriously fool different deep neural networks into making wrong predictions.

We further evaluate the classification performance of different deep models on the original test images in UCM and AID datasets and the adversarial images in the UAE-RS dataset. The overall accuracy $OA = n_{correct}/n_{total}$ is utilized as the evaluation metric, where $n_{correct}$ denotes the number of correctly classified samples. The detailed quantitative results are presented in Table VII. We find that most of the models reported here can achieve an OA of more than 90% in original clean test sets, which demonstrates the great capability of these

deep neural networks. Nevertheless, the performance of all models drops significantly on the adversarial test set. The OA gap between the clean set and adversarial set can reach more than 70 percentage points in many cases. These results indicate that adversarial examples in UAE-RS can easily mislead the existing state-of-the-art deep neural networks to make the wrong classification with high success rates on the scene classification of very high-resolution remote sensing images.

2) *UAE-RS Samples on Semantic Segmentation*: Some example images in the original Vaihingen and Zurich Summer datasets and the corresponding adversarial images in the UAE-RS dataset are presented in Figs. 8 and 9. We also visualize the segmentation maps of SegNet on both the original clean images and the adversarial images. It can be observed that although there appears to be little difference between the generated adversarial images and the original clean images for the human eye, adversarial images can seriously fool SegNet into making wrong predictions. While the segmentation maps of SegNet are very close to the ground-truth annotations, most pixels in the adversarial images are misclassified as the low vegetation category in the Vaihingen dataset or the building category in the Zurich Summer dataset.

We further evaluate the segmentation performance of dif-

ferent deep models on the original test images in Vaihingen and Zurich Summer datasets and the adversarial images in the UAE-RS dataset. The mean F_1 score $mF_1 = \frac{1}{n_v} \sum_{k=1}^{n_v} F_1^{(k)}$ is utilized as the evaluation metric, where $F_1^{(k)}$ denotes the F_1 score of the k th category and can be calculated as:

$$F_1^{(k)} = 2 \cdot \frac{\text{precision}^{(k)} \cdot \text{recall}^{(k)}}{\text{precision}^{(k)} + \text{recall}^{(k)}}. \quad (19)$$

The detailed quantitative results are presented in Table VIII. We find that most of the models reported here can achieve a mean F_1 score of more than 70% in original clean test sets. By contrast, the performance of all models drops significantly on the adversarial test set, especially for the Vaihingen dataset. The mean F_1 gap between the clean set and adversarial set can reach more than 30 percentage points in many cases. These results indicate that adversarial examples in UAE-RS can also fool most of the existing state-of-the-art deep neural networks on the semantic segmentation of very high-resolution remote sensing images. Besides, we also find networks that make use of global context information generally possess stronger resistance towards adversarial attacks. For example, the ContextNet shows relatively smaller mean F_1 gaps in both Vaihingen and Zurich Summer datasets. This could be a possible direction for researchers developing a robust deep model to address the threat of adversarial attacks in the future.

V. CONCLUSIONS AND DISCUSSIONS

In this study, we analyze the universal adversarial examples in remote sensing data for the first time. Specifically, we propose the Mixup-Attack and Mixcut-Attack methods to conduct black-box adversarial attacks. While most of the existing methods directly attack the predicted logits, the proposed methods aim to attack the shallow layer of deep neural networks by minimizing the KL-divergence between features of the virtual image (i.e., mixup or mixcut samples) and the input image. Despite their simplicity, extensive experiments on four benchmark very high-resolution remote sensing image datasets demonstrate that the proposed methods can generate transferable adversarial examples that cheat most of the state-of-the-art deep neural networks in both scene classification and semantic segmentation tasks with high success rates.

The experimental results in this study also indicate that the networks with deeper architectures generally possess stronger resistance against the black-box adversarial attack. Besides, the networks that can make use of global context information show stronger resistance towards adversarial attacks for the semantic segmentation task (e.g., ContextNet). These two aspects could also be potential directions for researchers developing a robust deep model to address the threat of adversarial attacks in the future. Another intriguing phenomenon is that Mixcut-Attack can achieve higher success rates on the attacks of scene classification tasks, while Mixup-Attack is more threatening to the attacks of semantic segmentation tasks according to our experimental results. One possible reason is that mixcut samples only perturb the real spatial distribution of different pixels in the clean image but cannot change the pixel values. Thus, mixcut samples could be helpful in attacking the scene

classification task where only one label is predicted for the whole image but would have less influence on the semantic segmentation task where pixel-wise classification is required.

We further collect the generated universal adversarial examples in the dataset named UAE-RS, which is the first dataset that provides black-box adversarial samples in the remote sensing field. We hope UAE-RS may serve as a benchmark for researchers developing adversarial defenses in the future. Despite the serious threat that adversarial examples have brought to deep learning models, it is also reported that training with adversarial examples could result in a better regularization ability and learn domain-invariant features [66]. Thus, there also exists some potential for researchers to address the domain adaptation problem in remote sensing tasks via adversarial examples, which can further be investigated as a possible future work.

ACKNOWLEDGMENT

The authors would like to thank Prof. Shawn Newsam for making the UCM dataset public available, Prof. Gui-Song Xia for providing the AID dataset, the International Society for Photogrammetry and Remote Sensing (ISPRS), and the German Society for Photogrammetry, Remote Sensing and Geoinformation (DGPF) for providing the Vaihingen dataset, and Dr. Michele Volpi for providing the Zurich Summer dataset.

REFERENCES

- [1] S.-A. Boukabara, J. Eyre, R. A. Anthes, K. Holmlund, K. M. S. Germain, and R. N. Hoffman, "The earth-observing satellite constellation: A review from a meteorological perspective of a complex, interconnected global system with extensive applications," *IEEE Geosci. Remote Sens. Mag.*, 2021.
- [2] P. Ghamisi, B. Rasti, N. Yokoya, Q. Wang, B. Hofle, L. Bruzzone, F. Bovolo, M. Chi, K. Anders, R. Gloaguen *et al.*, "Multisource and multitemporal data fusion in remote sensing: A comprehensive review of the state of the art," *IEEE Geosci. Remote Sens. Mag.*, vol. 7, no. 1, pp. 6–39, 2019.
- [3] Y. Xu, B. Du, L. Zhang, D. Cerra, M. Pato, E. Carmona, S. Prasad, N. Yokoya, R. Hansch, and B. L. Saux, "Advanced multi-sensor optical remote sensing for urban land use and land cover classification: Outcome of the 2018 ieee grss data fusion contest," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 6, pp. 1709–1724, 2019.
- [4] G. Cheng, J. Han, and X. Lu, "Remote sensing image scene classification: Benchmark and state of the art," *Proceedings of the IEEE*, vol. 105, no. 10, pp. 1865–1883, 2017.
- [5] X. Wu, W. Li, D. Hong, R. Tao, and Q. Du, "Deep learning for unmanned aerial vehicle-based object detection and tracking: A survey," *IEEE Geosci. Remote Sens. Mag.*, 2021.
- [6] O. Ghorbanzadeh, D. Tiede, L. Wendt, M. Sudmanns, and S. Lang, "Transferable instance segmentation of dwellings in a refugee camp-integrating cnn and obia," *European Journal of Remote Sensing*, vol. 54, no. sup1, pp. 127–140, 2021.
- [7] L. Zhang, L. Zhang, and B. Du, "Deep learning for remote sensing data: A technical tutorial on the state of the art," *IEEE Geosci. Remote Sens. Mag.*, vol. 4, no. 2, pp. 22–40, 2016.
- [8] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [9] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [10] N. Akhtar, A. Mian, N. Kardan, and M. Shah, "Advances in adversarial attacks and defenses in computer vision: A survey," *IEEE Access*, 2021.
- [11] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," *arXiv preprint arXiv:1611.01236*, 2016.

- [12] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*, 2017.
- [13] S. Komkov and A. Petiushko, "Advhat: Real-world adversarial attack on arcfac face id system," *arXiv preprint arXiv:1908.08705*, 2019.
- [14] Y. Xu, B. Du, and L. Zhang, "Assessing the threat of adversarial examples on deep neural networks for remote sensing scene classification: Attacks and defenses," *IEEE Trans. Geos. Remote Sens.*, vol. 59, no. 2, pp. 1604–1617, 2021.
- [15] L. Chen, G. Zhu, Q. Li, and H. Li, "Adversarial example in remote sensing image recognition," *arXiv preprint arXiv:1910.13222*, 2019.
- [16] A. Chan-Hon-Tong, G. Lenczner, and A. Plyer, "Demotivate adversarial defense in remote sensing," *arXiv preprint arXiv:2105.13902*, 2021.
- [17] W. Czaja, N. Fendley, M. Pekala, C. Ratto, and I.-J. Wang, "Adversarial examples in remote sensing," in *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2018, pp. 408–411.
- [18] L. Chen, Z. Xu, Q. Li, J. Peng, S. Wang, and H. Li, "An empirical study of adversarial examples on remote sensing image scene classification," *IEEE Trans. Geos. Remote Sens.*, vol. 59, no. 9, pp. 7419–7433, 2021.
- [19] Y. Xu, B. Du, and L. Zhang, "Self-attention context network: Addressing the threat of adversarial attacks for hyperspectral image classification," *IEEE Trans. Image Process.*, vol. 30, pp. 8671–8685, 2021.
- [20] N. Narodytska and S. P. Kasiviswanathan, "Simple black-box adversarial attacks on deep neural networks," in *CVPR Workshops*, vol. 2, 2017.
- [21] C. Guo, J. Gardner, Y. You, A. G. Wilson, and K. Weinberger, "Simple black-box adversarial attacks," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2484–2493.
- [22] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, "Understanding neural networks through deep visualization," *arXiv preprint arXiv:1506.06579*, 2015.
- [23] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii, "Virtual adversarial training: a regularization method for supervised and semi-supervised learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 8, pp. 1979–1993, 2018.
- [24] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *IEEE Symposium on Security and Privacy*, 2017, pp. 39–57.
- [25] H. Xu, Y. Ma, H.-C. Liu, D. Deb, H. Liu, J.-L. Tang, and A. K. Jain, "Adversarial attacks and defenses in images, graphs and text: A review," *International Journal of Automation and Computing*, vol. 17, no. 2, pp. 151–178, 2020.
- [26] V. Q. Vo, E. Abbasnejad, and D. C. Ranasinghe, "Query efficient decision based sparse attacks against black-box deep learning models," *arXiv preprint arXiv:2202.00091*, 2022.
- [27] Q. Zhang, X. Li, Y. Chen, J. Song, L. Gao, Y. He, and H. Xue, "Beyond imagenet attack: Towards crafting adversarial examples for black-box domains," *arXiv preprint arXiv:2201.11528*, 2022.
- [28] B. Ru, A. Cobb, A. Blaas, and Y. Gal, "Bayesopt adversarial attack," in *International Conference on Learning Representations*, 2019.
- [29] L. Meunier, J. Atif, and O. Teytaud, "Yet another but more efficient black-box adversarial attack: tiling and evolution strategies," *arXiv preprint arXiv:1910.02244*, 2019.
- [30] J. Du, H. Zhang, J. T. Zhou, Y. Yang, and J. Feng, "Query-efficient meta attack to deep neural networks," *arXiv preprint arXiv:1906.02398*, 2019.
- [31] Y. Li, S. Bai, Y. Zhou, C. Xie, Z. Zhang, and A. Yuille, "Learning transferable adversarial examples via ghost networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 11 458–11 465.
- [32] Y. Liu, X. Chen, C. Liu, and D. Song, "Delving into transferable adversarial examples and black-box attacks," *arXiv preprint arXiv:1611.02770*, 2016.
- [33] W. Zhou, X. Hou, Y. Chen, M. Tang, X. Huang, X. Gan, and Y. Yang, "Transferable adversarial perturbations," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 452–467.
- [34] Q. Huang, I. Katsman, H. He, Z. Gu, S. Belongie, and S.-N. Lim, "Enhancing adversarial example transferability with an intermediate level attack," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4733–4742.
- [35] S. Chen, Z. He, C. Sun, J. Yang, and X. Huang, "Universal adversarial attack on attention and the resulting dataset damagenet," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [36] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.
- [37] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo, "Cutmix: Regularization strategy to train strong classifiers with localizable features," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6023–6032.
- [38] H. Zhang, Y. Yu, J. Jiao, E. Xing, L. El Ghaoui, and M. Jordan, "Theoretically principled trade-off between robustness and accuracy," in *International Conference on Machine Learning*. PMLR, 2019, pp. 7472–7482.
- [39] L. Schwinn, R. Raab, A. Nguyen, D. Zanca, and B. Eskofier, "Exploring misclassifications of robust neural networks to enhance adversarial attacks," *arXiv preprint arXiv:2105.10304*, 2021.
- [40] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 9185–9193.
- [41] Y. Yang and S. Newsam, "Bag-of-visual-words and spatial extensions for land-use classification," in *Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 2010, pp. 270–279.
- [42] G.-S. Xia, J. Hu, F. Hu, B. Shi, X. Bai, Y. Zhong, L. Zhang, and X. Lu, "Aid: A benchmark data set for performance evaluation of aerial scene classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 7, pp. 3965–3981, 2017.
- [43] M. Cramer, "The dgpf-test on digital airborne camera evaluation overview and test design," *PFG Photogrammetrie, Fernerkundung, Geoinformation*, pp. 73–82, 2010.
- [44] M. Volpi and V. Ferrari, "Semantic segmentation of urban scenes by learning local class interactions," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, 2015, pp. 1–9.
- [45] Y. Hua, D. Marcos, L. Mou, X. X. Zhu, and D. Tuia, "Semantic segmentation of remote sensing images with sparse annotations," *IEEE Geosci. Remote Sens. Lett.*, 2021.
- [46] J. Lin, C. Song, K. He, L. Wang, and J. E. Hopcroft, "Nesterov accelerated gradient and scale invariance for adversarial attacks," *arXiv preprint arXiv:1908.06281*, 2019.
- [47] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Proc. Neural Inf. Process. Syst.*, vol. 32, pp. 8026–8037, 2019.
- [48] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Proc. Neural Inf. Process. Syst.*, vol. 25, pp. 1097–1105, 2012.
- [49] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [50] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4700–4708.
- [51] I. Radosavovic, R. P. Kosaraju, R. Girshick, K. He, and P. Dollár, "Designing network design spaces," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 10 428–10 436.
- [52] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3431–3440.
- [53] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 234–241.
- [54] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2881–2890.
- [55] A. Chaurasia and E. Culurciello, "Linknet: Exploiting encoder representations for efficient semantic segmentation," in *IEEE Visual Communications and Image Processing*. IEEE, 2017, pp. 1–4.
- [56] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [57] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2818–2826.
- [58] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1492–1500.
- [59] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, 2017.

- [60] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-decoder with atrous separable convolution for semantic image segmentation,” in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 801–818.
- [61] V. Badrinarayanan, A. Kendall, and R. Cipolla, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [62] H. Zhao, X. Qi, X. Shen, J. Shi, and J. Jia, “Icnnet for real-time semantic segmentation on high-resolution images,” in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 405–420.
- [63] R. P. Poudel, U. Bonde, S. Liwicki, and C. Zach, “Contextnet: Exploring context and detail for semantic segmentation in real-time,” *arXiv preprint arXiv:1805.04554*, 2018.
- [64] M. Treml, J. Arjona-Medina, T. Unterthiner, R. Durgesh, F. Friedmann, P. Schuberth, A. Mayr, M. Heusel, M. Hofmarcher, M. Widrich *et al.*, “Speeding up semantic segmentation for autonomous driving,” in *Proc. Neural Inf. Process. Syst. Workshops*, 2016.
- [65] T. Pohlen, A. Hermans, M. Mathias, and B. Leibe, “Full-resolution residual networks for semantic segmentation in street scenes,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4151–4160.
- [66] Y. Na, J. H. Kim, K. Lee, J. Park, J. Y. Hwang, and J. P. Choi, “Domain adaptive transfer attack-based segmentation networks for building extraction from aerial images,” *IEEE Trans. Geos. Remote Sens.*, vol. 59, no. 6, pp. 5171–5182, 2020.