

Using LIP to Gloss Over Faces in Single-Stage Face Detection Networks

Siqi Yang, Arnold Wiliem, Shaokang Chen, and Brian C. Lovell

The University of Queensland, Australia

siqi.yang@uq.net.au, a.wiliem@uq.edu.au, shaokangchenuq@gmail.com,
lovell@itee.uq.edu.au

Abstract. This work shows that it is possible to fool/attack recent state-of-the-art face detectors which are based on the single-stage networks. Successfully attacking face detectors could be a serious malware vulnerability when deploying a smart surveillance system utilizing face detectors. In addition, for the privacy concern, it helps prevent faces being harvested and stored in the server. We show that existing adversarial perturbation methods are not effective to perform such an attack, especially when there are multiple faces in the input image. This is because the adversarial perturbation specifically generated for one face may disrupt the adversarial perturbation for another face. In this paper, we call this problem the Instance Perturbation Interference (IPI) problem. This IPI problem is addressed by studying the relationship between the deep neural network receptive field and the adversarial perturbation. Besides the single-stage face detector, we find that the IPI problem also exists on the first stage of the Faster-RCNN, the commonly used two-stage object detector. As such, we propose the Localized Instance Perturbation (LIP) that confines the adversarial perturbation inside the Effective Receptive Field (ERF) of a target to perform the attack. Experimental results show the LIP method massively outperforms existing adversarial perturbation generation methods – often by a factor of 2 to 10.

Keywords: Adversarial · Interference · Effective receptive field · Single-stage network · Detection.

1 Introduction

Deep neural networks have achieved great success in recent years on many applications [31, 10, 5, 28, 15, 27, 6, 39, 17]. However, it has been demonstrated in various works that by adding tiny, imperceptible perturbations onto the image, the network output can be changed significantly [32, 4, 16, 25, 23, 11, 35, 19]. These perturbations are often referred to as adversarial perturbations [4]. Most prior works are primarily aimed at generating adversarial perturbations to fool neural networks for image classification tasks [32, 4, 16, 25, 23, 11, 22]. It is relatively easier to attack these networks as the perturbations need to change only one network decision for each image containing an instance/object of interest. This means, there is only a single target and the target is the entire image. Recently,




Fig. 1. An illustration of the Instance Perturbation Interference (IPI) problem. *Upper row:* two instances with their generated adversarial perturbations. The outer and inner circles indicate the Theoretical Receptive Field (TRF) and Effective Receptive Field (ERF), respectively. *Lower row:* one dimensional representation of the perturbations. IPI problem refers to the perturbation generated for one instance significantly disrupting the perturbation generated for the other instance. The disruption does not have significant effect on the left case, whereas on the right case, it will reduce the effectiveness of the attack

several methods have been proposed on more challenging attacks for segmentation [19, 3, 2] and object detection tasks [35], where there are significantly more targets to attack within the input image.

In the field of biometrics, Sharif *et al.* [29] showed that face recognition systems can be fooled by applying adversarial perturbations, where a detected face can be recognized as another individual. In addition, for the privacy concern, biometric data in a dataset might be utilized without the consent of the users. Therefore, Mirjalili *et al.* [20, 21] developed a technique to protect the soft biometric privacy (*e.g.*, gender) without harming the accuracy of face recognition. However, in the above-mentioned methods, the faces are still captured and stored in a server. In this paper, we propose a novel way to address these privacy issues by avoiding the faces be detected completely from an image. Thus, attacking face detection is crucial for both the security and privacy concerns.

With similar goals, previous works [29, 36] performed attacks on the Viola & Jones (VJ) face detector [33]. However, deep neural networks have been shown to be extremely effective in detecting faces [1, 6, 39, 12, 26, 37, 13, 24, 40], which can achieve 2 times higher detection rate than the VJ. In this work, we tackle the problem of generating effective adversarial perturbations for deep learning based face detection networks. To the best of our knowledge, this is the first study that attempts to perform such an adversarial attack on face detection networks.

Deep network based object/face detection methods can be grouped as two-stage network, *e.g.*, Faster-RCNN [28] and single-stage network [15, 27, 6, 24, 40]. In Faster-RCNN [9], a shallow region proposal network is applied to generate candidates and a deep classification network is utilized for the final decision. The Single-Stage (SS) network is similar to the region proposal network in Faster-RCNN [28] but performs both object classification and localization simultaneously. By utilizing the Single-Stage network architecture, recent detectors [6, 24, 40] can detect faces on various scales with a much faster running time. Due to

their excellent performance, we confine this paper to attacking the most recent face detectors utilizing Single-Stage network.

We find that applying the commonly-used gradient based adversarial methods [4, 23] to the state-of-the-art face detection networks has not presented satisfactory results. We point out that attacking a Single-Stage detector is challenging and the unsatisfactory performance is attributed to the Instance Perturbation Interference (IPI) problem. The IPI problem can be briefly explained as interference between the perturbation required to attack one instance and the perturbation required to attack a nearby instance. Since the recent adversarial perturbation methods [35, 19] do not consider this problem, they become quite ineffective in attacking SS face detector networks.

In this work, we attribute the IPI problem to the receptive field of deep neural networks. Recent work [18] shows that the receptive field follows a 2D Gaussian distribution, where the set of input image pixels closer to an output neuron have higher impact on the neuron decision. The area where high impact pixels are concentrated is referred to as the Effective Receptive Field (ERF) [18]. As illustrated in Fig. 1, if two faces are close to each other, the perturbation generated to attack one face will reside in the ERF of another face. Prior work [34] shows that adversarial attacks might fail when the specific structure is destroyed. Thus, the residency in the ERF significantly hampers the success of attacking the other face. In other words, the IPI problem happens when the interfering perturbations disrupt the adversarial perturbations generated for the neighboring faces. This IPI problem will become more serious when multiple faces exist in close proximity and when the receptive field of the network is large. For the general two-stage object detection, Faster-RCNN [28], we find that the IPI problem also exists on its first stage network, *i.e.*, region proposal network (RPN). We believe this is the first work that describes and explains the IPI problem.

Contributions - We list our contributions as follows: (1) We describe and provide theoretical explanation of the Instance Perturbation Interference problem that makes the existing adversarial perturbation generation method fail to attack the SS face detector networks when multiple faces exist; (2) This is the first study to show that it is possible to attack deep neural network based face detector. More specifically, we propose an approach to attack Single-Stage based face detector networks. (3) To perform the attack, we propose Localised Instance Perturbation (LIP) method to generate instance based perturbations by confining the perturbations inside each instance ERF.

2 Background

2.1 Adversarial Perturbation

As mentioned, attacking a network means attempting to change the network decision on a particular target. A target t is defined as a region in the input image where the generated adversarial perturbation is added to change the network decision corresponding to this region. For example, the target t for attacking an image classification network is the entire image.

The adversarial perturbation concept was first introduced for attacking image classification networks in [32, 4, 16, 25, 23, 11, 22]. Szegedy *et al.* [32] showed that by adding imperceptible perturbations to the input images, one could make the Convolutional Neural Network (CNN) predict the wrong class label with high confidence. Goodfellow *et al.* [4] explained that the vulnerability of the neural networks to the adversarial perturbations is caused by the linear nature of the neural networks. They proposed a fast method to generate such adversarial perturbations, naming the method Fast Gradient Sign Method (FGSM) defined by: $\boldsymbol{\xi} = \alpha \text{sign}(\nabla_{\mathbf{X}} \ell(f(\mathbf{X}), y^{true}))$, where α was a hyper-parameter [4]. The gradient was computed with respect to the entire input image $\mathbf{X} \in \mathbb{R}^{w \times h}$ by back-propagation and the function $\text{sign}()$ is the L_∞ norm. Following this, Kurakin *et al.* [11] proposed to extend FGSM by iteratively generating the adversarial perturbations. At each iteration, the values of the perturbations were clipped to control perceptibility. We denote it as I-FGSM in this work. To reduce perceptibility, Moosavi-Dezfooli *et al.* [23] proposed the method DeepFool, which iteratively adds the minimal adversarial perturbations to the images by assuming the classifier was linear at each iteration. The existence of universal perturbations for image classification was shown in [22].

More recently, adversarial examples were extended into various applications such as semantic segmentation [35, 19, 2, 3] and object detection [35]. Metzen *et al.* [19] adapted the I-FGSM described in [11] into the semantic segmentation domain, where every pixel was a target. They demonstrated that the gradients of the loss for different target pixels might point to the opposite directions. In object detection, the instances of interest are the detected objects. Thus, the targets are the detected region proposals containing the object. An approach for generating adversarial perturbations for object detection is proposed in [35]. They claimed that generating adversarial perturbations in object detection was more difficult than in the semantic segmentation task. In order to successfully attack a detected object, one needs to ensure all the region proposals associated with the object/instance are successfully attacked. For example, if only K out of R region proposals are successfully attacked, the detector can still detect the object by using the other high-confidence-score region proposals that are not successfully attacked.

We note that all of the above approaches use whole image perturbations which have the same size as the input image. This is because these perturbations are generated by calculating the gradient with respect to the entire image. Thus, a generated perturbation for one target may disrupt the perturbations generated for other targets. To contrast these methods with our work, we categorize these methods as **IMage based Perturbation (IMP)** methods.

2.2 Loss function

In general, the perturbations are generated by optimizing a specific objective function. Let $\mathcal{L} = \sum_{i=1}^T \mathcal{L}_{t_i}$ be the loss function to optimize. The objective function is defined as follows:

$$\arg \min_{\boldsymbol{\xi}} \sum_{i=1}^T \mathcal{L}_{t_i}(\boldsymbol{\xi}), \quad (1)$$

where T is the number of targets; \mathcal{L}_{t_i} is the loss function for each individual target t_i ; and $\boldsymbol{\xi} \in \mathbb{R}^{w \times h}$ is the adversarial perturbation which will be added into the input image \mathbf{X} .

According to the goals of adversarial attacks, the attacks can be categorized into *non-targeted adversarial attacks* [4, 22, 35] and *targeted adversarial attacks* [11, 19]. For non-targeted adversarial attacks, the goal is to reduce the probability of truth class y^{true} of the given target t and to make the network predict any arbitrary class, whereas the goal of targeted adversarial attacks is to ensure the network predict the target class y^{target} for the target t . The objective function of the targeted attacks can be summarized into the following formula:

$$\arg \min_{\boldsymbol{\xi}} \mathcal{L}_{t_i} = \ell(f(\mathbf{X} + \boldsymbol{\xi}, t_i), y^{target}) - \ell(f(\mathbf{X} + \boldsymbol{\xi}, t_i), y^{true}), \quad (2)$$

where, $\boldsymbol{\xi}$ is the optimum adversarial perturbation; f is the network classification score matrix on the target region; and ℓ is the network loss function.

In general, the face detection problem is considered as a binary classification problem, which aims at classifying a region as face (+1) or non-face (-1) (*i.e.*, $y^{target} = \{+1, -1\}$). However, in order to detect faces in various scales, especially for tiny faces, recent face detectors utilizing Single-Stage networks [6, 24, 40] divide the face detection problem into multiple scale-specific binary classification problems, and learn their loss functions jointly. The objective function to attack such a network is defined as:

$$\arg \min_{\boldsymbol{\xi}} \mathcal{L}_{t_i} = \sum_{j=1}^S \ell_{s_j}(f_{s_j}(\mathbf{X} + \boldsymbol{\xi}, t_i), y^{target}), \quad (3)$$

where, S is the number of scales; and ℓ_{s_j} is the scale-specific detector loss function. Compared to Eq. 2, the above objective is more challenging. This is because a single face can not only be detected by multiple region proposals/targets, but also by multiple scale-specific detectors. Thus, one can only successfully attack a face when the adversarial perturbation fools all the scale-specific detectors. In other words, attacking the single-stage face detection network is more challenging than the work in object detection [35].

Finally, as our main aim is to prevent faces being detected, then our objective function is formally defined as:

$$\mathcal{L} = \sum_{i=1}^T \mathcal{L}_{t_i} = \sum_{i=1}^T \sum_{j=1}^S \ell_{s_j}(f_{s_j}(\mathbf{X} + \boldsymbol{\xi}, t_i), -1). \quad (4)$$

In this work, we use the recent state-of-the-art Single-Stage face detector, HR [6], which jointly learns 25 different scale-specific detectors, *i.e.*, $S = 25$.

3 Instance Perturbations Interference

When performing an attack using the existing adversarial perturbation approaches [11, 19], the Instance Perturbations Interference (IPI) problem appears when multiple faces exist in the input image. In short, the IPI problem refers to the conditions where successfully attacking one instance of interest can reduce the chance of attacking the other instances of interest. For the face detection task, the instance of interest is a face. If not addressed, the IPI problem will

significantly reduce the overall attack success rate. To show the existence of the IPI problem, we perform an experiment using synthetic images. In this experiment, we apply an adaptation of the existing perturbation methods generated by minimizing Eq. 4.

3.1 Image based Perturbation

As mentioned, we categorize the previous methods as Image based Perturbation (IMP) as they use whole image perturbation to perform the attack. Here we adapt two of the existing methods, I-FSGM [11] and DeepFool [23], by optimizing Eq. 4. We denote them as IMP(I-FGSM) and IMP(DeepFool). In both methods, the adversarial perturbation is generated by using a gradient descent approach. At the $(n + 1)$ th iteration, the gradient with respect to the input image \mathbf{X} , $\nabla_{\mathbf{X}}\mathcal{L}(f(\mathbf{X} + \boldsymbol{\xi}^{(n)}), -1)$, is generated via back-propagating the network with the loss function.

For the IMP(I-FSGM) [11], we iteratively update the adversarial perturbation as follows:

$$\boldsymbol{\xi}^{(n+1)} = \text{Clip}_{\varepsilon}\{\boldsymbol{\xi}^{(n)} - \alpha \text{sign}(\nabla_{\mathbf{X}}\mathcal{L}(f(\mathbf{X} + \boldsymbol{\xi}^{(n)}), -1))\}, \quad (5)$$

where the step rate $\alpha = 1$; the epsilon ε is the maximum absolute magnitude to clip; $\boldsymbol{\xi}^{(0)} = \mathbf{0}$; and the loss function \mathcal{L} is referred to the Eq. 4. Note that in Eq. 4, the loss function is a summation of the loss of all targets. Thus, the aggregate gradient, $\nabla_{\mathbf{X}}\mathcal{L}$, can be rewritten as:

$$\nabla_{\mathbf{X}}\mathcal{L}(f(\mathbf{X} + \boldsymbol{\xi}^{(n)}), -1) = \sum_{i=1}^T \sum_{j=1}^S \nabla_{\mathbf{X}}\ell_{s_j}(f_{s_j}(\mathbf{X} + \boldsymbol{\xi}^{(n)}, t_i), -1). \quad (6)$$

As we assume f is a deep neural network, then the aggregate gradient $\nabla_{\mathbf{X}}\mathcal{L}$ can be obtained by back-propagating all of the targets at once. After obtaining the final adversarial perturbation $\boldsymbol{\xi}$, the perturbed image, \mathbf{X}^{adv} , is then generated by: $\mathbf{X}^{adv} = \mathbf{X} + \boldsymbol{\xi}$.

For the IMP(DeepFool), following [23], we configure the Eq. 5 into:

$$\boldsymbol{\xi}^{(n+1)} = \text{Clip}_{\varepsilon}\left\{\boldsymbol{\xi}^{(n)} - \frac{\nabla_{\mathbf{X}}\mathcal{L}(f(\mathbf{X} + \boldsymbol{\xi}^{(n)}))}{\|\nabla_{\mathbf{X}}\mathcal{L}(f(\mathbf{X} + \boldsymbol{\xi}^{(n)}))\|_2}\right\}, \quad (7)$$

where the loss function in Eq. 4 is rewritten as $\mathcal{L} = \sum_{i=1}^T \sum_{j=1}^S (f_{s_j}(\mathbf{X} + \boldsymbol{\xi}, t_i))$.

Compare with the IMP(DeepFool), the IMP(I-FGSM) generates denser and more perceptible perturbations due to the L_{∞} norm.

3.2 Existence of the IPI Problem

To show the existence of the IPI problem, we construct a set of synthetic images by controlling the number of faces and distances between them: (1) an image containing only one face; (2) an image containing multiple faces closely located in a grid and (3) using image in (2) but increasing the distance between the faces. Examples are shown in Fig. 2. For this experiment, we use the recent state-of-the-art face detector HR-ResNet101 [6]. The synthetic images are constructed by randomly selecting 50 faces from the WIDER FACE dataset [38]. Experimental details are given in Section 5.2. We generate the adversarial perturbations using the IMP approaches: IMP(I-FGSM) and IMP(DeepFool).

The attack success rate is calculated as follows: $\frac{\#\text{Face removed}}{\#\text{Detected face}}$. Table 1 reports the results. For the first synthetic case where an image only contains one face, both IMP(I-FGSM) and IMP(DeepFool) are able to attack the face detector with a 100% attack success rate. The IMP method is only partially successful on the second case where the number of faces is increased to 16. The attack success rates decrease significantly to only 18.3% and 11.0% when $N = 81$. The IMP method attack success rates significantly increase when the distances between faces are increased significantly, especially for the IMP(DeepFool). It is because the IMP(DeepFool) generates sparser perturbations than the IMP(I-FGSM).

These results suggest the following: (1) IMP is effective when only a single face exists; (2) IMP is ineffective when multiple faces exist close to each other and (3) the distance between faces significantly affects the attack performance. There are two questions that arise from these results: (1) why is the attack affected by the number of faces? and (2) why does the distance between faces affect the attack success rate? We address these two questions in the next section.




Fig. 2. Examples of synthetic images after adding adversarial perturbations from the Image based Perturbation (IMP). The detection results of the adversarial images are shown in rectangles. Note that, as face density increases, attack success rate decreases. The IMP attack is ineffective when there are many faces in an image, as in (a) and (c). When the distance among faces is increased, the attack becomes successful as in (b)

4 Proposed method

We first elaborate on the relationship between the Effective Receptive Field and the IPI problem. Then, the proposed Localized Instance Perturbation (LIP) method is outlined.

4.1 Effective Receptive Field (ERF)

The receptive field of a neuron in a neural network is a set of pixels in the input image that impact the neuron decision [18]. In CNNs, it has been shown in [18] that the distribution of impact within the Theoretical Receptive Field (TRF) of a neuron follows a 2D Gaussian distribution. This means most pixels that have significant impact to the neuron decision are concentrated near the neuron and the impact decays quickly away from the center of the TRF. In [18], the area where pixels still have significant impact to the neuron decision is defined as the Effective Receptive Field (ERF). The ERF only takes up a fraction of the TRF and pixels within the ERF will generate non-negligible impacts on the final outputs. We argue that understanding ERF and TRF is important for addressing the IPI problem. This is because the adversarial perturbation is aimed at changing a network decision at one or more neurons. All pixels in the input image that impact the decision must be considered.

Table 1. The IMP attack success rate (in %) on the synthetic images with respect to the number of faces and distances among faces. N is the number of faces. The IMP can achieve 100% attack success rate when there is one face per image. The attack success rate drops significantly when the number of faces is increased. With the same number of faces, the attack success rate can be increased as the distance among faces increases

N	1	16		81	
Distance	40	40	80	160	40
IMP(I-FGSM)	100	34	37.5	38.9	18.3

(a) IMP(I-FGSM)

N	1	16			81
Distance	40	40	60	80	40
IMP(DeepFool)	100	67.5	91.8	99.7	11.0

(b) IMP(DeepFool)

In this paper, we denote the Distribution of Impacts in the TRF as DI-TRF for simplicity. The DI-TRF is measured by calculating the partial derivative of the central pixel on the output layer via back-propagation. Following the notations in our paper, let us denote the central pixel as t_c , then the partial derivative of the central pixel is $\frac{\partial f(\mathbf{X}, t_c)}{\partial \mathbf{X}}$, which is the DI-TRF. According to the chain rule, we have the gradient of the target t_c [18] as: $\nabla_{\mathbf{X}} \mathcal{L}(f(\mathbf{X}, t_c), y^{target}) = \frac{\partial \mathcal{L}(f(\mathbf{X}, t_c), y^{target})}{\partial f(\mathbf{X}, t_c)} \frac{\partial f(\mathbf{X}, t_c)}{\partial \mathbf{X}}$, where the $\frac{\partial \mathcal{L}(f(\mathbf{X}, t_c), y^{target})}{\partial f(\mathbf{X}, t_c)}$ is set to 1.

Comparing the gradient of a target pixel for the adversarial perturbations in Eq. 6, the only difference with the DI-TRF is in the partial derivative of the loss function $\frac{\partial \mathcal{L}(f(\mathbf{X}, t_c), y^{target})}{\partial f(\mathbf{X}, t_c)}$, which is a scalar for one target pixel. In our work, the scalar, $\frac{\partial \mathcal{L}(f(\mathbf{X}, t_c), y^{target})}{\partial f(\mathbf{X}, t_c)}$, measures the loss between the prediction label and the target label. The logistic loss is used for the binary classification of each scale-specific detector, (*i.e.*, the $\ell_{s_j}(f_{s_j}(\mathbf{X}, t_c), y^{target})$ in Eq. 4). Therefore, our adversarial perturbation for one target can be considered as a scaled distribution of the DI-TRF. Since DI-TRF follows a 2D Gaussian distribution [18], then the adversarial perturbation to change a single neuron decision is also a 2D Gaussian.

We explain the IPI problem as follows. Since an adversarial perturbation to attack a single neuron follows a 2D Gaussian, then the perturbation is mainly spread over the ERF and will have a non-zero tail outside the ERF. From the experiment, we observed that the perturbations generated to attack multiple faces in the image may interfere with other. More specifically, when these perturbations overlap with the neighboring face ERF, they may be sufficient enough to disrupt the adversarial perturbation generated to attack this neighboring face. In addition, prior work [34] shows that adversarial attacks might fail when the specific structure is destroyed. In other words, when multiple attacks are applied simultaneously, these attacks may corrupt each other, leading to a lower attack rate. We name the part of a perturbation interfering with the other perturbations for other faces as the interfering perturbation.

This also explains why the IPI is affected by the distance between faces. The closer the faces, the more chance the interfering perturbations with a larger magnitude overlap with the neighboring face ERF. When distances between faces increase, the magnitude of the interfering perturbations that overlap with the neighboring ERFs may not be strong enough to disrupt attacks for target faces.

4.2 Localized Instance Perturbation (LIP)

To address the IPI problem, we argue that the generated adversarial perturbations of one instance should be exclusively confined within the instance ERF. As such, we call our method as the Localized Instance Perturbation (LIP). The LIP comprises two main components: (1) methods to eliminate any possible interfering perturbation and (2) methods to generate the perturbation.

Eliminating the Interfering Perturbation. To eliminate the interference between perturbations, we attempt to constrain the generated perturbation for each instance individually inside the ERF. Let us consider that an image \mathbf{X} , with $w \times h$ pixels, contains N instances $\{\mathbf{m}_i\}_{i=1}^N$. Each instance \mathbf{m}_i has its corresponding ERF, \mathbf{e}_i , and we have $\{\mathbf{e}_i\}_{i=1}^N$. For each instance, there are a set of corresponding targets represented as object proposals, $\{p_j\}_{j=1}^P$. We denote the final perturbation for the i th instance as \mathbf{R}_{m_i} and the final combination of the perturbation of all the instances as \mathbf{R} . Similar to the IMP method, once the final perturbation, \mathbf{R} , has been computed, then we add the perturbation into the image $\mathbf{X}^{adv} = \mathbf{X} + \mathbf{R}$.

(1) *Perturbation Cropping.* This step is to limit the perturbations inside the instance ERF. This is done by cropping the perturbation according to the corresponding instance ERF. Let us define a binary matrix $\mathbf{C}_{\mathbf{e}_i} \in \{0, 1\}^{w \times h}$ as the cropping matrix for the ERF, \mathbf{e}_i . The matrix \mathbf{C} is defined as follows:

$$\mathbf{C}_{\mathbf{e}_i}(w, h) = \begin{cases} 1, & (w, h) \in \mathbf{e}_i \\ 0, & \text{otherwise} \end{cases}, \quad (8)$$

where (w, h) is a pixel location. The cropping operation is computed by a element-wise dot product of the mask $\mathbf{C}_{\mathbf{e}_i}$ and the gradient w.r.t. the input images \mathbf{X} , is defined as:

$$\mathbf{R}_{m_i} = \mathbf{C}_{\mathbf{e}_i} \cdot \nabla_{\mathbf{X}} \mathcal{L}_{m_i}, \quad (9)$$

where \mathcal{L}_{m_i} is the loss function of the i -th instance. \mathcal{L}_{m_i} will be described in the next sub-section.

(2) *Individual Instance Perturbation.* It is possible to compute the perturbation of multiple instances simultaneously. However, the interfering perturbation can still exist and may impact the attack. To that end, we propose to separately compute the perturbation for each instance, $\nabla_{\mathbf{X}} \mathcal{L}_{m_i}$ before cropping. After the cropping step is applied to each instance perturbation, the final perturbation of all instances is combined via:

$$\mathbf{R} = \sum_{i=1}^N \mathbf{C}_{\mathbf{e}_i} \cdot \nabla_{\mathbf{X}} \mathcal{L}_{m_i}. \quad (10)$$

We then normalize the final perturbation, \mathbf{R} , via: $\mathbf{R} = \alpha \text{sign}(\mathbf{R})$.

Perturbation Generation. Given a set of region proposals corresponding to an instance, there are at least two methods of generating the instance perturbation \mathbf{R}_{m_i} : (1) All proposal based generation and (2) Highest Confidence proposal based generation.

(1) *All Proposal based Generation.* In the first method, we utilize all the region proposals to generate the perturbation \mathbf{R}_{m_i} . Thus, the \mathcal{L}_{m_i} in Eq. 9 can be defined as a summation of the loss function of all the region proposals \mathcal{L}_{p_j} belong to the instance:

$$\mathcal{L}_{m_i} = \sum_{j=1}^P \mathcal{L}_{p_j} . \quad (11)$$

(2) *Highest Confidence Proposal based Generation.* In online hard example mining [30], Shrivastava *et al.* showed the efficiency of using the hard examples to generate the gradients for updating the networks. The hard examples are the high-loss object proposals chosen by the non-maximum suppression. Non-Maximum Suppression (NMS) is similar to max-pooling, which selects the object proposal with the highest score (*i.e.*, selecting the proposal with the highest loss). Inspired by this, instead of attacking all of the object proposals corresponding to a single instance, we can use NMS to select the one with the highest loss to compute the back-propagation. Then \mathcal{L}_{m_i} can be rewritten as:

$$\mathcal{L}_{m_i} = \max(\mathcal{L}_{p_j}) . \quad (12)$$

5 Experiments

5.1 Implementation Details

In this section, we first describe the implementation details and then evaluate our proposed adversarial attacks on the state-of-the-art face detection datasets.

For this study, we utilize a recent state-of-the-art face detector, HR [6]. In particular, HR-ResNet101 is used. Image pyramids are utilized in HR, *i.e.*, down-sampling/interpolating the input image into multiple sizes. Therefore, for every image in the pyramid, we generate corresponding adversarial examples. The detection results of the image pyramid are combined together with Non-Maximum Suppression (NMS). The chosen thresholds of NMS and classification are 0.1 and 0.5 respectively.

In order to avoid the gradient explosion when generating the perturbations, we found that by zero-padding the small input images can reduce the magnitude of the gradient. In this work, we zero pad the small images to 1000×1000 pixels. In addition, as the input images of the detection networks can have arbitrary sizes, we do not follow existing methods [22, 19] that resize the input images into a canonical size.

Note that we cannot simply crop the input image to generate a successful adversarial perturbation. This is because the perturbation may be incomplete as it does not include the context information obtained from neighboring instances. An example of two non-normalized perturbations in absolute value generated with and without context is shown in supplementary materials.

For determining the perturbation cropping size, we follow the work of Luo *et al.* [18] which computes the gradient of the central proposal of an instance on the output feature map to obtain the distribution of the ERF. We average the gradients over multiple instances and determine the crop size with the definition that the ERF takes up 90% of the energy of the TRF [18]. The perturbation

crop size is set to 80×80 pixels for small faces and 140×140 pixels for large faces. The maximum noise value ε is 20 and the maximum number of iterations N_0 is 40. The α is set to 1 in this work.

Perturbation Generation Methods. In our work, we compared our proposed Localized Instance Perturbation (LIP) approach with the Image Perturbation (IMP) and Localized Perturbation (LP). The details of the perturbation generation methods evaluated are listed as follow:

(1) *Localized Instance Perturbation using All proposal generation (LIP-A)*. The proposed LIP-A is a variant of our proposed LIP method in Section 4.2. As mentioned, the loss function of one instance is the summation of all proposals (refer to Eq. 11).

(2) *Localized Instance Perturbation using Highest confidence proposal generation (LIP-H)*. The LIP-H is another variant of our proposed LIP with the loss function of Eq. 12. The loss function of one instance consists of only one loss of the highest confidence proposal.

(3) *Image Perturbation (IMP)*. The IMP method refers to the generation method in Section 3.1 which applies the perturbation without cropping it. This perturbation generation method follows the previous work [19].

(4) *Localized Perturbation (LP)*. The LP is the localized perturbation which also crops the image perturbation. The main difference to the proposed LIP is that it computes the gradients of all the instances simultaneously before the cropping. In contrast to Eq. 10, the final perturbation is obtained by:

$$\mathbf{R} = \bigcup_{i=1}^N \mathbf{C}_{e_i} \cdot \sum_{i=1}^N \nabla_{\mathbf{X}} \mathcal{L}_{m_i} . \quad (13)$$

where $\bigcup_{i=1}^N \mathbf{C}_{e_i}$ is the union of all binary matrices. The advantage of this method is that current deep learning toolboxes can calculate the summation of the gradients of all instances, (*i.e.*, $\sum_{i=1}^N \nabla_{\mathbf{X}} \mathcal{L}_{m_i}$), simultaneously by back-propagating the network only once.

Benchmark Datasets. We evaluate our proposed adversarial perturbations on two recent popular face detection benchmark datasets: (1) *FDDDB dataset [8]*: The FDDDB dataset includes images of faces with a wide range of difficulties such as occlusions, difficult poses, low resolution and out-of-focus faces. It contains 2,845 images with a total of 5,171 faces labeled; and (2) *WIDER FACE dataset [38]*: The WIDER FACE dataset is currently the most challenging face detection benchmark dataset. It comprises 32,203 images and 393,703 annotated faces based on 61 events collected from the Internet. The images of some events, *e.g.*, parade, contain a large number of faces. According to the difficulties of the occlusions, poses, and scales, the faces are grouped into three sets: 'Easy', 'Medium' and 'Hard'.

Evaluation Metrics. The metrics for evaluating the adversarial attacks against face detection are defined as follows: (1) *Attack Success Rate*: The attack success rate is the ratio between the number of faces that are successfully attacked and the number of detected faces before the attacks; and (2) *Detection Rate*: The detection rate is the ratio between the number of detected faces and the number of faces in the images.

5.2 Evaluation on Synthetic Data

As discussed in Section 3, due to the IPI problem, the IMP does not perform well on the cases where (1) the number of faces per image is large; and (2) the faces are close to each other. Here, we contrast IMP with LP and LIP.

We randomly selected 50 faces from the WIDER FACE dataset [38]. These faces were first resized into a canonical size of 30×30 pixels. Each face was then duplicated and inserted into a blank image in a rectangular grid manner (*e.g.*, $3 \times 3 = 9$). The number of duplicates and the distance between the duplicates were controlled during the experiment. In total there were 50 images and the attack success rate was then averaged across 50 images. Some examples of the synthetic images are shown in Fig. 2.

The Effect of the Number of Faces. We progressively increased the number of duplication for each synthetic image from 1×1 to $9 \times 9 = 81$ duplicates. We fixed the distance between duplicates to 40 pixels. The quantitative results are shown in Fig. 3. From this figure, we can see that for the perturbation generation method I-FGSM, the IMP attack success rate significantly drops from 100% to 20% as the number of faces is increased. On the contrary, both LP and LIP-H can achieve significantly higher attack success rate than IMP. This is because both LP and LIP-H only use the generated perturbation within the corresponding instance ERF by cropping it before applying. Note that, when the number of faces is more than 36, the LP attack success rate drops from 85% ($N = 36$) to 51% ($N = 81$), whereas the LIP-H can still achieve more than 90% success rate. Since LP processes all the instances simultaneously, the accumulation of the interfering perturbations within each instance ERF will become more significant when the number of faces is increased. Similarly, for the generation method DeepFool, the LIP has demonstrated its effectiveness on addressing the IPI problem when multiple faces exist. These also suggest the existence of the IPI problem.

The Effect of Distance between Faces. In this experiment the number of faces duplication was fixed to 9. We modified the distance between face duplicates to 40, 160 and 240 pixels. It can be seen from Fig. 3b that the attack success rate for IMP increases as the distance between faces is increased. The performance of both LP and LIP-H are not affected. Similar performance is achieved on the DeepFool. More details are shown in the supplementary materials.




Fig. 3. The attack success rate of the I-FGSM with respect to: (a) the number of faces. The distance was fixed to 40 pixels; and (b) the distance between faces. Nine face duplicates were used. (c) The attack success rate of the DeepFool

5.3 Evaluation on Face Detection Datasets

We contrasted LIP-A and LIP-H with IMP and LP based on two existing methods: I-FGSM [11] and DeepFool [23]. The experiments were run on the FDDB [8] and 1,000 randomly selected images in the WIDER FACE validation set [38].

The results based on the I-FGSM, are reported in Table 3 and 2 respectively. On the FDDB dataset (in Table 3), the face detector, HR [6], achieves 95.7% detection rate. The LP, LIP-A and LIP-H can significantly reduce the detection rate to around 5% with the attack success rate of 94.9%, 94.6% and 93.8% respectively. On the other hand, the IMP can only achieve 53.9% attack success rate (*i.e.*, significantly lower than the LP, LIP-A, LIP-H performance). This signifies the importance of the perturbation cropping to eliminate the interfering perturbations. Due to the IPI problem, the interfering perturbations from the other instances will affect the adversarial attacks of the target instance. This results in the low attack success rate of the IMP. This is because to generate the perturbations, the IMP simply sums up the all perturbations including the interfering perturbations. We note that the performance of LP, LIP-A and LIP-H are on par in the FDDB dataset. This could be due to the low number of faces per image for this dataset.

However, when the number of faces per image increases significantly, LIP shows its advantages. Examples can be seen in Fig. 4. This can be observed in the WIDER FACE dataset (in Table 2) where LIP-A and LIP-H outperform LP by 4 percentage points. the LIP-H can achieve attack success rates of (69.8%, 63.7%, 61.4%)




Fig. 4. Examples of the adversarial attacks on face detection network, where the perturbation generation is based on the I-FGSM. The LIP-H is successfully attack all the faces, whereas some faces are still detected when IMP is used

Table 3. The attack success and detection rates (in %) on FDDB [8]

Perturbations	none	I-FGSM			
		IMP	LP	LIP-A	LIP-H
Detection Rate	95.7	44.1	4.8	5.1	5.9
Attack Success Rate	–	53.9	94.9	94.6	93.8

Table 2. The attack success and detection rate (in %) on WIDER FACE [38]

Perturbations		none	I-FGSM				DeepFool	
			IMP	LP	LIP-A	LIP-H	IMP	LIP-A
Detection Rate	easy	92.4	46.2	30.1	28.2	26.5	50.6	43.2
	medium	90.7	50.7	34.7	32.2	31.1	54.4	40.0
	hard	77.3	45.9	29.3	23.6	26.6	46.5	25.8
Attack Success Rate	easy	–	50.0	67.4	69.5	71.3	45.3	53.2
	medium	–	44.1	61.7	64.5	65.7	40.0	56.4
	hard	–	40.6	62.1	69.5	65.6	39.6	66.6

on the (easy, medium, hard) sets, while the LP can only obtain attack success rate (65.7%, 59.5%, 57.4%). As the LP processes all the instances together, the interfering perturbations are accumulated within the ERF before the cropping step. Note that the interfering perturbations may have low magnitude, however, when they are accumulated due to the number of neighboring instances then disruption could be significant. These results also suggest that we do not necessarily need to attack all the region proposals as the performance of LIP-H is on par with LIP-A. Similarly, for the DeepFool based methods, the LIP has demonstrated its effectiveness on addressing the IPI problem.

5.4 Evaluation on Object Detection Dataset

To explore the existence of the IPI problem in object detection networks, we perform attacks on the pre-trained Faster-RCNN [28] (based on ResNet101 [5]) provided by the Tensorflow object detection API [7]. More specifically, we attack the 1st stage (*i.e.*, RPN) of Faster-RCNN with the goal of reducing generated proposals. We choose 300 images from COCO2017 dataset [14], where the average number of objects per image is 15. The original predicted detections from the pre-trained Faster-RCNN are taken as ground truth. The results in Table 4 show that the IPI problem exists and our proposed LP method can attack more than 60% of the instances that cannot be attacked by IMP. Note that, as the RPN generates hundreds of proposals for each instance, the proposed LIP methods are not used due to the high computations.

Table 4. Evaluation on COCO2017 dataset [14]

Perturbations	IMP	LP
Average Recall	7.9	2.2
Average Precision	6.9	1.9

6 Conclusions

In this paper, we presented an adversarial perturbation method to fool a recent state-of-the-art face detector utilizing the single-stage network. We described and addressed the Instance Perturbation Interference (IPI) problem which was the root cause for the failure of the existing adversarial perturbation generation methods to attack multiple faces simultaneously. We found that it was sufficient to only use the generated perturbations within an instance/face Effective Receptive Field (ERF) to perform an effective attack. In addition, it was important to exclude perturbations outside the ERF to avoid disrupting other instance perturbations. We thus proposed the Localized Instance Perturbation (LIP) approach that only confined the perturbation within the ERF. Experiments showed that the proposed LIP successfully generated perturbations for multiple faces simultaneously to fool the face detection network and outperformed existing adversarial generation methods. In the future, we plan to develop a universal perturbation generation method which can attack many faces with a general perturbation.

Acknowledgements This work has been funded by Sullivan Nicolaides Pathology, Australia, and the Australian Research Council (ARC) Linkage Projects Grant LP160101797. Arnold Wiliem is funded by the Advance Queensland Early-Career Research Fellowship.

References

1. Chen, D., Hua, G., Wen, F., Sun, J.: Supervised transformer network for efficient face detection. In: European Conference on Computer Vision (ECCV). Springer (2016)
2. Cisse, M., Adi, Y., Neverova, N., Keshet, J.: Houdini: Fooling deep structured prediction models. In: Advances in Neural Information Processing Systems (NIPS) (2017)
3. Fischer, V., Kumar, M.C., Metzen, J.H., Brox, T.: Adversarial examples for semantic image segmentation. In: International Conference on Learning Representations (ICLR) workshop (2017)
4. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: International Conference on Learning Representations (ICLR) (2015)
5. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2016)
6. Hu, P., Ramanan, D.: Finding tiny faces. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2017)
7. Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S., et al.: Speed/accuracy trade-offs for modern convolutional object detectors. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2017)
8. Jain, V., Learned-Miller, E.G.: Fddb: A benchmark for face detection in unconstrained settings. UMass Amherst Technical Report (2010)
9. Jiang, H., Learned-Miller, E.: Face detection with the faster r-cnn. In: IEEE International Conference on Automatic Face & Gesture Recognition (FG). IEEE (2017)
10. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems (NIPS) (2012)
11. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. In: International Conference on Learning Representations (ICLR) workshop (2017)
12. Li, H., Lin, Z., Shen, X., Brandt, J., Hua, G.: A convolutional neural network cascade for face detection. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2015)
13. Li, Y., Sun, B., Wu, T., Wang, Y.: Face detection with end-to-end integration of a convnet and a 3d model. In: European Conference on Computer Vision (ECCV). Springer (2016)
14. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European Conference on Computer Vision (ECCV). pp. 740–755. Springer (2014)
15. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: European Conference on Computer Vision (ECCV). Springer (2016)
16. Liu, Y., Chen, X., Liu, C., Song, D.: Delving into transferable adversarial examples and black-box attacks. In: International Conference on Learning Representations (ICLR) (2017)
17. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2015)
18. Luo, W., Li, Y., Urtasun, R., Zemel, R.: Understanding the effective receptive field in deep convolutional neural networks. In: Advances in Neural Information Processing Systems (NIPS) (2016)

19. Metzen, J.H., Kumar, M.C., Brox, T., Fischer, V.: Universal adversarial perturbations against semantic image segmentation. In: International Conference on Computer Vision (ICCV). IEEE (2017)
20. Mirjalili, V., Raschka, S., Namboodiri, A., Ross, A.: Semi-adversarial networks: Convolutional autoencoders for imparting privacy to face images. In: International Conference on Biometrics (ICB) (2018)
21. Mirjalili, V., Ross, A.: Soft biometric privacy: Retaining biometric utility of face images while perturbing gender. In: International Joint Conference on Biometrics (IJCB) (2017)
22. Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2017)
23. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2016)
24. Najibi, M., Samangouei, P., Chellappa, R., Davis, L.: Ssh: Single stage headless face detector. In: International Conference on Computer Vision (ICCV). IEEE (2017)
25. Nguyen, A., Yosinski, J., Clune, J.: Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2015)
26. Qin, H., Yan, J., Li, X., Hu, X.: Joint training of cascaded cnn for face detection. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2016)
27. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2016)
28. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems (NIPS) (2015)
29. Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K.: Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. ACM (2016)
30. Shrivastava, A., Gupta, A., Girshick, R.: Training region-based object detectors with online hard example mining. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2016)
31. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. International Conference on Learning Representations (ICLR) (2015)
32. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. In: International Conference on Learning Representations (ICLR) (2014)
33. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2001)
34. Xie, C., Wang, J., Zhang, Z., Ren, Z., Yuille, A.: Mitigating adversarial effects through randomization. In: International Conference on Learning Representations (ICLR) (2018)
35. Xie, C., Wang, J., Zhang, Z., Zhou, Y., Xie, L., Yuille, A.: Adversarial examples for semantic segmentation and object detection. In: International Conference on Computer Vision (ICCV). IEEE (2017)

36. Yamada, T., Gohshi, S., Echizen, I.: Privacy visor: Method for preventing face image detection by using differences in human and device sensitivity. In: IFIP International Conference on Communications and Multimedia Security. Springer (2013)
37. Yang, S., Luo, P., Loy, C.C., Tang, X.: From facial parts responses to face detection: A deep learning approach. In: International Conference on Computer Vision (ICCV) (2015)
38. Yang, S., Luo, P., Loy, C.C., Tang, X.: Wider face: A face detection benchmark. In: Computer Vision and Pattern Recognition (CVPR). IEEE (2015)
39. Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters* **23**(10), 1499–1503 (2016)
40. Zhang, S., Zhu, X., Lei, Z., Shi, H., Wang, X., Li, S.Z.: S3 fd: Single shot scale-invariant face detector. In: International Conference on Computer Vision (ICCV). IEEE (2017)