

Generative Dynamic Patch Attack

Xiang Li
xli62@gsu.edu
Shihao Ji
sji@gsu.edu

Department of Computer Science
Georgia State University
Atlanta, GA, USA

Abstract

Adversarial patch attack is a family of attack algorithms that perturb a part of image to fool a deep neural network model. Existing patch attacks mostly consider injecting adversarial patches at input-agnostic locations: either a predefined location or a random location. This attack setup may be sufficient for attack but has considerable limitations when using it for adversarial training. Thus, robust models trained with existing patch attacks cannot effectively defend other adversarial attacks. In this paper, we first propose an end-to-end patch attack algorithm, Generative Dynamic Patch Attack (GDPA), which generates both patch pattern and patch location adversarially for each input image. We show that GDPA is a generic attack framework that can produce dynamic/static and visible/invisible patches with a few configuration changes. Secondly, GDPA can be readily integrated for adversarial training to improve model robustness to various adversarial attacks. Extensive experiments on VGGFace, Traffic Sign and ImageNet show that GDPA achieves higher attack success rates than state-of-the-art patch attacks, while adversarially trained model with GDPA demonstrates superior robustness to adversarial patch attacks than competing methods. Our source code can be found at <https://github.com/lxuniverse/gdpa>.

1 Introduction



Figure 1: Different types of patch attacks: (a) Eyeglasses Attack [28], (b) Sticker Attack [8], (c) Adversarial Patch [8], and (d) GDPA (ours).

Deep neural networks (DNNs) have demonstrated remarkable success in solving complex prediction tasks in a variety of fields: computer vision [8], natural language processing [30] and speech recognition [26]. However, recent studies show that they are particularly vulnerable to adversarial examples [12] in the form of small perturbations to inputs that lead DNNs to predict incorrect outputs.

Recent works [8, 8, 17, 28, 36, 37], show that perturbing part of an image with perceivable noise is another effective method to attack neural network models. Typically, attackers

can craft perceivable patches to replace part of images for adversarial attack. The advantage of this perceivable patch attack is that it is more practical than the imperceptible adversarial attacks in the real world: adversaries can paste a sticker on a traffic sign to attack the autopilot system of autonomous vehicles. There are several situations where patch attack is significant concerning due to its security threats: 1) an attacker uses adversarially designed eyeglass frames [28] to fool face recognition (Fig. 1a), 2) an attacker pastes adversarially crafted stickers [8] on stop signs to fool traffic sign classification (Fig. 1b), and 3) a universal adversarial patch [9] causes targeted misclassification of any object (Fig. 1c).

However, it is a significant limitation that most patch attack algorithms do not consider the problem of finding the best location in an image to inject the patch. Existing patch attack algorithms either use a fixed position as patch location [8, 28, 57] or learn patches that are universal across different locations [9, 17, 56]. The fixed location methods show high attack success rates but are poorly performed at other locations, while the random location patches do not have competitive attack success rates compared to the fixed location methods. To address this issue, in this paper we propose a Generative Dynamic Patch Attack (GDPA), which learns image-dependent patch pattern and patch location altogether. GDPA is inspired by the idea that different images have different sets of weak pixels since DNN classifiers typically focus on different image regions when queried by different images [29]. Therefore, an image-dependent dynamic patch attack would be more effective than a fixed location or random location patch attack.

On the other hand, due to the security threats of adversarial attacks, a variety of adversarial defense algorithms have been developed recently [12, 21, 22], among which adversarial training (AT) [12] has been proved the most effective one for hardening neural networks against adversarial attacks. Although AT with the PGD attack [22] is the most scalable and effective method for learning robust models, a recent work of Wu et al. [53] shows that AT exhibits limited effectiveness against three high-profile physically realizable patch attacks: eyeglasses attack [28], sticker attack [8] and adversarial patch [9]. To overcome this limitation, Wu et al. [53] propose a Rectangular Occlusion Attack (ROA) for adversarial training, which yields models highly robust to patch attacks. ROA is a two-stage patch attack algorithm, which first uses a *gray pattern* to find the location in image that maximizes the cross-entropy loss via grid search, and then optimizes the patch pattern at the identified position. However, this two-stage patch attack method is suboptimal and has quite a few limitations (see a discussion in Sec. 2), which motivates us to propose GDPA that learns patch pattern and patch location simultaneously. Moreover, to improve the inference efficiency, GDPA employs a generator to generate patch pattern and location with one forward propagation, without expensive iterative optimizations that are employed by other attack algorithms, such as PGD [22] and ROA [53]. Concretely, we make the following contributions:

- We introduce a generic patch attack method GDPA that can generate dynamic/static and visible/invisible patch attacks with a few configuration changes.
- GDPA employs a generator to generate patch pattern and patch location altogether per image, and reduces the inference time substantially (e.g., 40-50x faster).
- GDPA is an end-to-end differentiable patch attack algorithm and can be readily integrated for adversarial training to defend against high-profile patch attacks.
- Experiments show that GDPA has superior attack success rates over strong patch attack baselines, and the adversarially trained model with GDPA is more robust to various adversarial attacks than state-of-the-art methods.

2 Related Works

Adversarial Attack Most adversarial attack methods focus on adding imperceptible perturbation covering the entire image [11, 12, 61]. Recently, researchers have shown that perturbing a part of image with perceptible noise is another practical method to attack DNN models [3, 8, 17, 18, 19, 28, 33, 36, 37]. Sharif et al. [28] propose to add eyeglasses with a specially constructed frame texture to attack face recognition. Eykholt et al. [8] show that adding specific rectangular solid-colored patches on traffic signs can fool traffic sign classification. LAVAN [17] learns visible and localized patches that are transferable across images and locations by training the pattern at a random location with a randomly picked image in each iteration. Recently, Wu et al. [33] propose a Rectangle Occlusion Attack (ROA) to generate adversarial patches for adversarial training. ROA uses an exhaustive search (ROA-Exh) or a gradient guided search (ROA-Grad) to find the location that maximizes the cross-entropy (CE) loss and optimizes the patch pattern afterwards. Specifically, ROA-Exh exhaustively searches on images with a stride, and ROA-Grad uses the magnitude of gradient of the CE loss as the sensitivity of regions to identify the top candidate regions to accelerate the location search. However, ROA has some considerable limitations. Firstly, it employs a two-stage attack generation, which separates the process of finding the patch location and patch pattern into two steps: it first finds the position using a *gray pattern* and then optimizes the patch pattern at that position. Hence, the location identified by a *gray pattern* may not be the best patch location for the optimized pattern. Secondly, the two-stage optimization of ROA is computationally expensive and slows down the patch generation process during inference. Different from these algorithms, our GDPA trains a generator to generate the patch pattern and location altogether for each input image. Moreover, GDPA is end to end differentiable, which entails an efficient optimization and easy integration for adversarial training.

Before GDPA, several works [1, 24, 25, 34] have proposed to train generators to generate perturbation to improve the fooling rate and inference speed. Poursaeed et al. [24] present a trainable network to transform input images to adversarial perturbations. Baluja and Fischer [1] train feed-forward neural networks in a self-supervised manner to generate adversarial examples against a target network. Different to these generator-based attack methods, our GDPA generates both patch pattern and patch location altogether, and employ an affine transform to synthesize adversarial patch examples.

Adversarial Defense Defending against adversarial attacks is a challenging task. Different types of defense algorithms have been proposed in the past few years [11, 9, 7, 9, 13, 14, 14, 15, 20, 21, 23, 33, 35], among which adversarial training (AT) [22] has been proved the most effective one against adversarial attacks. AT employs adversarial examples as data augmentation to train a robust model. It has been shown that this method can improve the defense accuracy effectively and sometimes can even improve the accuracy upon the model trained only on the original clean dataset [32]. However, a recent work of Wu et al. [33] shows that robust models trained by AT exhibit limited effectiveness against high-profile patch attacks [3, 8, 28]. As the first work attempting to defend patch attacks, Wu et al. [33] propose DOA, which performs a standard adversarial training with Rectangle Occlusion Attack (ROA). As we discussed earlier in this section, ROA has some considerable limitations, which limit its performance on adversarial defense. Our GDPA does not suffer from those limitations of ROA, and is end-to-end differentiable and more amenable for adversarial training.

3 The GDPA Framework




Figure 2: The GDPA generation pipeline. Given an image x , GDPA generates a patch pattern and a patch location for weighted adversarial patch injection. $\alpha \in [0, 1]$ controls the visibility of the patch attack. The pipeline is fully differentiable.

GDPA is a framework that aims to conduct dynamic patch attack by generating adversarial patch pattern and patch location altogether for each input image. It has a generic formulation that can generate dynamic/static and visible/invisible patch attacks. As an overview, Figure 2 illustrates the GDPA generation pipeline, while Figure 3 demonstrates how GDPA can be utilized to train an adversarially robust model.

3.1 Problem Formulation

We start with the definition of dynamic patch attack. Let $\mathbb{D} = \{\mathcal{X}, \mathcal{Y}\}$ denote a training dataset, where \mathcal{X} is a set of images of size $w \times h$, and \mathcal{Y} are their corresponding labels. Let $T : \mathcal{X} \rightarrow \mathcal{Y}$ denote a target model that we attempt to attack. Given an image $x \in \mathcal{X}$ and a target model T , our *dynamic patch attack* aims to find a pattern of size $w' \times h'$ and a position in image that once placed on image x it can mislead the target model.

3.2 Localized Pattern Generation

One crucial component of GDPA is the generator that generates patch pattern and patch location for a given image. Since patch pattern and patch location are coupled to a given image, we design a generator G with two heads that share the same latent features extracted by an encoder. Specifically, our generator includes an encoder G_E to extract the feature representation of image x , followed by a location decoder G_L and a pattern decoder G_P to generate location and pattern of the adversarial patch:

$$l_x, l_y = \tanh(G_L(G_E(x))/\beta), \quad (1)$$

$$pattern = 0.5 \times \tanh(G_P(G_E(x))) + 0.5, \quad (2)$$

where l_x and l_y are the location (2D coordinates) of a patch in image x with the origin at the center of image, and *pattern* is the patch pattern of size $w' \times h'$. To keep the patch location l_x and l_y within the boundary of image, we use a tanh function to constrain l_x and l_y in the range of $[-1, 1]$, where β is a hyperparameter that controls the slope of tanh. All experiments in this paper use $\beta = 3000$, which we found to work well across a variety of architectures and datasets. Similarly, we use another tanh to impose the pattern values in the range of $[0, 1]$ ¹.

Specially, we use a convolutional neural network as our encoder network G_E , with an architecture adapted from the work of image-to-image translation [69]. On top of G_E , we use two fully-connected networks as our decoders G_P and G_L , respectively. Due to page limit, details of the network architectures are provided in the Appendix.

¹As a preprocessing step, all images are normalized to have pixel values in the range of $[0, 1]$.

3.3 Weighted Adversarial Patch Injection

With the generated patch location and patch pattern, we then define a function to inject the patch into image x . Standard adversarial attacks [22] employ an additive function to inject noise: $x' = x + p$, where p is an imperceptible adversarial perturbation. Recently, other forms of perturbations, such as multiplicative ones $x' = x \odot m$ [58], have been explored to inject perturbations. In addition, LAVAN [17] employs $(1 - m) \odot x + m \odot p$ with a binary mask $m \in \{0, 1\}^{w' \times h'}$ to generate patch attack adversarial examples. Inspired by LAVAN, we extend this function by relaxing the binary mask to a continuous mask $m \in [0, 1]^{w' \times h'}$ for adversarial patch injection. Specifically, we employ the weighted adversarial patch injection $x^{adv} = (1 - m) \odot x + m \odot p$, with $m \in [0, 1]^{w' \times h'}$, which is a convex combination of original image x and patch pattern p with the weight defined by m . We find this relaxed version is more flexible and easier to optimize than the one LAVAN explored. Next, we discuss how to use the generated (l_x, l_y) and *pattern* to inject an adversarial patch to image x .

3.4 Differentiable Affine Transformation

We employ an affine transformation in GDPA to inject adversarial patches into images. To make the whole pipeline differentiable w.r.t. l_x and l_y , a bilinear interpolation is used to estimate the pixel values that are not on the pixel grids after transformation. By doing this, the whole pipeline is fully differentiable and the gradient can be back-propagated end-to-end to update parameters of generator G . Specifically, we adopt the affine transformation and image sampling method of Spatial Transformer Networks [16] to define a differentiable translate operator, which can translate a source image to a target image by a displacement of (l_x, l_y) .

We first use an affine transform to compute the pixel index relationship between source image and target image:

$$\begin{pmatrix} x_i^s \\ y_i^s \end{pmatrix} = \begin{bmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \theta_{21} & \theta_{22} & \theta_{23} \end{bmatrix} \begin{pmatrix} x_i^t \\ y_i^t \\ 1 \end{pmatrix}, \quad (3)$$

where (x_i^t, y_i^t) is the pixel index of target image, and (x_i^s, y_i^s) is the corresponding pixel index in source image. We set $\theta_{11} = 1, \theta_{21} = 0, \theta_{12} = 0, \theta_{22} = 1, \theta_{13} = w'/2 \cdot l_x$ and $\theta_{23} = h'/2 \cdot l_y$ for translation purpose². Thus, we have $x_i^s = x_i^t + w'/2 \cdot l_x$ and $y_i^s = y_i^t + h'/2 \cdot l_y$, where $l_x, l_y \in [-1, 1]$. Since (x_i^s, y_i^s) are continuous variables, we can use a bilinear interpolation to sample the pixel values from source image:

$$v_i = \sum_{j=0}^{w-1} \sum_{k=0}^{h-1} u_{jk} \max(0, 1 - |x_i^s - j|) \max(0, 1 - |y_i^s - k|), \quad (4)$$

where u_{jk} is the pixel value at index (j, k) of the source image, and v_i is the output value of pixel i at index (x_i^t, y_i^t) of the translated image. With the affine transform and bilinear sampler described above, we have a differentiable translate operator, which we denote as $Translate()$ in the rest part of the paper.

3.5 Generative Dynamic Patch Attack

Figure 2 illustrates the GDPA generation pipeline, which includes the three components we described above: patch pattern and location generator, differentiable affine transform, and the weighted adversarial patch injection to produce an image-dependent dynamic patch attack.

²Note that we can also learn $\theta_{11}, \theta_{12}, \theta_{21}, \theta_{22}$ to rotate, dilate or shear an adversarial patch to further improve GDPA's performance. For simplicity and also because we can already achieve state-of-the-art attack success rate (ASR) with translation, in this paper we only consider translation and will leave advanced affine transform to future work.




Figure 3: The GDPA-AT pipeline. Given an image, GDPA generates an adversarial patch to maximize the loss of classifier T , while classifier T learns from the patch attack to minimize its loss.

As shown in Figure 2, we introduce an initial mask m_{center} of the same size of input image, and the center part of the mask has value 1 and rest of 0. Then we use the affine transform $Translate()$ to translate m_{center} by a displacement of (l_x, l_y) :

$$m = \alpha \cdot Translate(m_{center}, l_x, l_y), \quad (5)$$

where $\alpha \in [0, 1]$ is a hyperparameter that controls the visibility of adversarial patches. When $\alpha = 1$, the patch would be completely visible and replace the original image pixel values; otherwise, the visibility of the adversarial patch will be lower. In practice, we can use a small value of α to generate human imperceptible adversarial patches.

Similarly, we can generate a translated patch pattern. As shown in Figure 2, once *pattern* is generated, we zero-pad it to create a pattern p_{center} of the same size of input image with *pattern* at the center. We then translate p_{center} by (l_x, l_y) via the affine transform: $p = Translate(p_{center}, l_x, l_y)$.

Finally, we can generate a GDPA adversarial example for image x by $x^{adv} = (1 - m) \odot x + m \odot p$. As we can see, all the components in Figure 2 are differentiable. Therefore, the whole GDPA generation pipeline is fully differentiable and can be optimized efficiently with gradient-based methods.

$$\operatorname{argmin}_G -\mathcal{L}_{CE}(T, x^{adv}, y). \quad (6)$$

We can also launch a targeted patch attack to fool the target model T to misclassify an input x as target class

$$\operatorname{argmin}_G \mathcal{L}_{CE}(T, x^{adv}, y_{target}). \quad (7)$$

Details of the GDPA training algorithm can be found in Algorithm 1.

3.6 Adversarial Training with GDPA

Adversarial training with the PGD attack exhibits limited effectiveness against high-profile patch attacks [33]. In this section, we discuss how to utilize GDPA for adversarial training to improve model robustness against high-profile patch attacks.

Figure 3 illustrates the GDPA adversarial training (GDPA-AT) pipeline to train a robust model against patch attacks. Similar to Generative Adversarial Networks [14], GDPA-AT trains generator G and target classifier T iteratively to optimize the following minimax objective:

$$\min_T \max_G \mathbb{E}_{(x,y) \sim \mathbb{D}} [\mathcal{L}_{CE}(T, x^{adv}, y)], \quad (8)$$

where the inner maximization step optimizes generator G to maximize the classification loss of T , while the outer minimization step optimizes target classifier T to minimize the classification loss. Unlike the traditional adversarial training, in which the inner maximization step

usually optimizes an adversarial example x^{adv} directly, our GDPA-AT optimizes a generator G to generate patch attack with one forward propagation. As the iterative training proceeds, the generator G searches for the weakest image region to attack classifier T at each iteration, while T learns from the current patch attacks and becomes more resilient to these attacks over time. Details of our GDPA-AT algorithm are described in Algorithm 2.

Algorithm 1: GDPA Training

Input: training set \mathbb{D} ; target model T ;
visibility α
Output: generator G
initialize generator G ;
for number of training epochs **do**
 for each $(x, y) \in \mathbb{D}$ **do**
 $l_x, l_y = G_L((G_E(x)));$
 $pattern = G_P(G_E(x));$
 $m = \alpha \cdot Translate(m_{center}, l_x, l_y);$
 $p = Translate(p_{center}, l_x, l_y);$
 $x^{adv} = (1 - m) \odot x + m \odot p;$
 if targeted attack **then**
 $loss = \mathcal{L}_{CE}(T, x^{adv}, y_{target});$
 else
 $loss = -\mathcal{L}_{CE}(T, x^{adv}, y);$
 end
 $\theta_G = \theta_G - lr * \partial loss / \partial \theta_G$
 end
end

Algorithm 2: GDPA-AT Training

Input: training set \mathbb{D}
Output: target classifier T ; generator G
initialize classifier T and generator G ;
for number of training epochs **do**
 for each $(x, y) \in \mathbb{D}$ **do**
 $x^{adv} = GDPA(G, x);$
 $loss = -\mathcal{L}_{CE}(T, x^{adv}, y);$
 $\theta_G = \theta_G - lr_G * \partial loss / \partial \theta_G;$
 $x^{adv} = GDPA(G, x);$
 $loss = \mathcal{L}_{CE}(T, x^{adv}, y);$
 $\theta_T = \theta_T - lr_T * \partial loss / \partial \theta_T$
 end
end

4 Experimental Results

We now validate GDPA on benchmark datasets for adversarial patch attack and adversarial defense. Specifically, we evaluate the performance of GDPA on patch attack in Section 4.1 and GDPA-AT on improving model robustness in Section 4.2. To evaluate the inference efficiency, we also compare the run-times of GPDA and state-of-the-art attack algorithms in Section 4.3. All our experiments are performed with PyTorch on Nvidia RTX GPUs. Our source code is provided as a part of supplementary materials.

Experimental Setup We evaluate GDPA and GDPA-AT on three benchmark datasets: VGGFace [23], Traffic Sign [8] and ImageNet [6]. To evaluate GDPA’s attack performance, we compare GDPA with LAVAN [17] and ROA [63], two state-of-the-art patch attack algorithms that generate patches based on iterative optimizations. Following their experimental settings, we run LAVAN and ROA for 50 optimization iterations with a learning rate of 4. For adversarial defense experiments, we compare GDPA-AT with DOA [63] and PGD-AT [22]. The former is a state-of-the-art defense algorithm for patch attacks, while the latter is a well-established defense algorithm for adversarial attacks. We evaluate the robustness of the models under eyeglasses attack [23], sticker attack [8]. Following the settings in DOA [63], we use 70×70 patches with stride 5 for VGGFace and 7×7 patches with stride 2 for Traffic Sign to generate ROA attacks. We set $\epsilon = 16$ for PGD-AT since this yields the best results of PGD-AT. We use attack success rate (ASR) [6] as the metric to evaluate the effectiveness of an attack, and use classification accuracy to evaluate the robustness of a model when under adversarial attacks. Details of benchmark datasets, high-profile patch attacks, network architectures and training procedures can be found in the Appendix.

4.1 Dynamic Patch Attack

We first evaluate the performance of GDPA on non-targeted and targeted patch attacks and compare it with the state-of-the-arts: LAVAN [17] and ROA [13]. We provide results of two versions of ROA: ROA-Exh and ROA-Grad, where the former exhaustively searches for a patch location in images with a fixed stride, and the latter uses the magnitude of gradient as the sensitivity of regions to identify top regions to accelerate the location search. We evaluate the effectiveness of the attack algorithms when perturbing different percentages of pixels. To interpret the results, we also visualize the perturbed images generated by GDPA.

Dataset	Algorithm	Percentage of Attacked Pixels							
		Non-Targeted Attack				Targeted Attack			
		1%	2%	5%	10%	1%	2%	5%	10%
Traffic Sign	LAVAN [17]	33.4	58.7	85.1	93.9	32.2	48.1	80.9	89.9
	ROA-Grad [13]	36.2	61.8	87.3	93.6	29.8	44.6	74.5	90.5
	ROA-Exh [13]	37.1	63.0	89.4	93.8	31.3	45.9	76.2	91.7
	GDPA	39.6	64.1	91.3	94.3	33.9	50.4	77.5	92.8
VGGFace	LAVAN [17]	31.9	42.7	56.3	92.0	37.8	57.9	67.2	94.6
	ROA-Grad [13]	37.5	62.3	84.2	99.6	46.3	75.6	89.0	99.2
	ROA-Exh [13]	38.3	64.5	86.0	99.6	48.2	76.7	91.1	99.3
	GDPA	46.3	76.4	88.4	99.5	50.5	83.4	95.5	99.8
ImageNet	LAVAN [17]	89.2	92.8	97.8	99.9	86.3	93.8	99.7	99.8
	ROA-Grad [13]	93.5	94.6	98.7	99.7	79.6	88.3	97.5	99.8
	ROA-Exh [13]	94.8	95.3	99.2	99.7	81.1	89.6	98.4	99.8
	GDPA	96.3	96.9	99.7	99.8	89.3	94.4	99.6	99.9

Table 1: The ASRs of different patch attack algorithms on datasets Traffic Sign, VGGFace and ImageNet. Both non-targeted attack and targeted attack are considered. The performances are evaluated with patches of different sizes.

Table 1 reports the ASRs of GDPA and the other competing algorithms for non-targeted and targeted patch attacks. The ASRs of an attack algorithm are evaluated on a model trained with cross-entropy (CE) loss when attacked with patches of different sizes (1%, 2%, 5% or 10% of pixels). Specifically, We use square patches of width 3, 5, 7, 10 for Traffic Sign and 23, 32, 50, 71 for VGGFace and ImageNet. For targeted attacks, we choose the first class for each of the three datasets as the target class, i.e., “AddedLine”, “Aamir Khan” and “tench, Tinca tinca”, respectively. As expected, the larger patch size is, the higher ASR is achieved for all patch attack algorithms. In most of the cases, GDPA achieves higher ASRs than the competing algorithms.



Figure 4: Perturbed images of VGGFace and ImageNet generated by GDPA with different patch sizes. The last column of targeted attack are example images of target classes.




Figure 5: The impact of α to the ASR of GDPA on VGGFace.

Dataset	Patch Location	Percentage of Attacked Pixels							
		Non-Targeted Attack				Targeted Attack			
		1%	2%	5%	10%	1%	2%	5%	10%
Traffic Sign	Random	35.1	59.4	86.3	93.5	28.9	44.1	75.6	90.1
	Dynamic	39.6	64.1	91.3	94.3	33.9	50.4	77.5	92.8
VGGFace	Random	35.4	58.8	82.7	96.5	43.9	71.7	88.3	97.9
	Dynamic	46.3	76.4	88.4	99.5	50.5	83.4	95.5	99.8
ImageNet	Random	91.6	93.5	98.1	99.7	79.3	87.0	96.6	99.8
	Dynamic	96.3	96.9	99.7	99.8	89.3	94.4	99.6	99.9

Table 2: The ASRs of GDPA with patches of dynamic or random locations on datasets Traffic Sign, VGGFace and ImageNet.

Figure 4 visualizes the perturbed images of different patch sizes generated by GDPA for non-targeted and targeted attack on VGGFace and ImageNet. As we can see, the patches generated on VGGFace (top row) demonstrates clear semantic meanings, resembling human eyes, nose or mouth. Moreover, the positions chosen by GDPA on face images are in a close proximity of the original face features. On the other hand, the patches generated on ImageNet (bottom row) do not demonstrate a strong semantic meaning with non-targeted attack but shows semantic meaning with targeted attack.

Visibility α vs. ASR We further investigate the impact of visibility parameter α of Eq. 5 to GDPA’s ASR. The results on VGGFace are shown in Figure 5, where we consider different patch sizes. As expected, when α increases, the attack strength of GDPA gets stronger for all different patch sizes. Notably, when the patch size is 5% or 10% of pixels, GDPA can reach almost the highest ASRs when $\alpha \geq 0.6$, indicating that when patches are sufficiently large, the attack can be more invisible to attack a model successfully. Example perturbed images generated by GDPA with different α are provided in the Appendix.

Dynamic vs. Random Patch Location Tab 2 shows the comparison of GDPA framework with random patch locations or dynamic patch locations. As we can see from the table, GDPA achieves higher ASR with dynamic patch location than random location. This shows that learned image-dependant dynamic locations contribute to the superior performance of GDPA.

Attack Strength	Eyeglasses Attack (VGGFace)				Sticker Attack (Traffic Sign)				LAVAN (ImageNet)			
	0	100	200	300	0	10	100	1000	0%	1%	2%	5%
CE Training	98.9	0.2	0.1	0.0	98.7	42.9	32.5	24.3	76.6	41.4	0.7	0.1
PGD-AT [14]	97.3	37.7	36.9	36.6	97.5	57.6	45.1	42.5	62.4	45.4	37.8	12.0
DOA-Grad [6]	99.2	85.3	83.7	81.9	95.6	85.2	84.8	82.8	61.9	52.8	51.1	47.3
DOA-Exh [6]	99.0	89.8	87.8	85.9	92.9	92.2	91.8	90.8	62.0	54.2	53.7	50.6
GDPA-AT	99.6	96.5	94.7	94.9	98.5	96.2	95.8	94.6	62.2	58.3	56.9	53.2

Table 3: The accuracies of different robust models under eyeglasses attack on VGGFace (left), sticker attack on Traffic Sign (middle), and LAVAN on ImageNet (right). We evaluate under different iterations for Eyeglasses Attack and Sticker Attack, and different patch size for LAVAN.

4.2 Dynamic Patch Adversarial Training

Next we validate the robustness of models trained by GDPA-AT against various adversarial patch attacks. Specifically, we report the results of GDPA-AT trained models against Eye-glass Attack, Sticker Attack and LAVAN, and compare them with state-of-the-art defense methods.

Table 3 reports the accuracies of robust models trained by different defense algorithms against three types of patch attacks: 1) eyeglasses attack on VGGFace, 2) sticker attack on Traffic Sign and 3) LAVAN on ImageNet. As can be seen, PGD-AT, a well-established defense method for conventional adversarial attacks, is not robust to all three patch attacks, which is consistent with the results reported in [63]. While both DOA and GDPA-AT improve the robustness over PGD-AT significantly, GDPA-AT achieves substantially higher accuracies than the two variants of DOA.

4.3 Inference Speed

Besides the improved attack and defense performance of GDPA, another advantage of GDPA is its superior inference speed to generate attacks over the optimization-based methods, such as PGD [22] and ROA [63]. To have a quantitative comparison in terms of inference time, we evaluate the run-time of GDPA, PGD and ROA on the VGGFace test dataset (470 images). GDPA needs one forward propagation to generate a patch attack, while we follow the settings of ROA and PGD and run 50 iterative optimizations to generate their attacks. As shown in Table 4, GDPA is about 40x faster than PGD and 47x faster than ROA.

	PGD [22]	ROA-Grad [63]	ROA-Ex [63]	GDPA
Inference-time (s)	108.31	129.42	248.82	2.76

Table 4: Inference-time comparison of different attack algorithms on the VGGFace test dataset (470 images).

Additional Experimental Results As GDPA is a generic attack algorithm, we conduct additional experiments to evaluate its performance with different configurations and also validate some design choices. Due to page limit, details are relegated to the Appendix.

5 Conclusion

This paper introduces GDPA, a novel *dynamic* patch attack algorithm, that generates patch pattern and patch location altogether for each input image. Due to its generic formulation, GDPA can generate dynamic/static and visible/invisible patch attacks. GDPA is end-to-end differentiable, which entails an efficient optimization and easy integration for adversarial training. We validated our method on multiple benchmarks with different model architectures. GDPA demonstrates superior ASR over strong patch attack methods, and the adversarially trained model with GDPA is more robust to high-profile patch attacks. Moreover, GDPA is 40-50x faster than competing attack algorithms, making it a highly effective attack and defense algorithm.

6 Acknowledgment

We would like to thank the anonymous reviewers for their comments and suggestions, which helped improve the quality of this paper. We would also gratefully acknowledge the support of Cisco Systems Inc. for its research fund to this work.

References

- [1] Naveed Akhtar, Jian Liu, and Ajmal Mian. Defense against universal adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [2] Shumeet Baluja and Ian Fischer. Adversarial transformation networks: Learning to generate adversarial examples. *arXiv preprint arXiv:1703.09387*, 2017.
- [3] Tom B Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017.
- [4] Ping-yeh Chiang, Renkun Ni, Ahmed Abdelkader, Chen Zhu, Christoph Studer, and Tom Goldstein. Certified defenses for adversarial patches. In *ICLR 2020*, 2020.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 2009.
- [6] Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. Benchmarking adversarial robustness on image classification. In *Conference on Computer Vision and Pattern Recognition*, 2020.
- [7] Gintare Karolina Dziugaite, Zoubin Ghahramani, and Daniel M Roy. A study of the effect of jpg compression on adversarial images. *arXiv preprint arXiv:1608.00853*, 2016.
- [8] Ivan Evtimov, Kevin Eykholt, Earlene Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust physical-world attacks on machine learning models. In *CoRR*, 2017.
- [9] Ji Gao, Beilun Wang, Zeming Lin, Weilin Xu, and Yanjun Qi. Deepcloak: Masking deep neural network models for robustness against adversarial samples. *arXiv preprint arXiv:1702.06763*, 2017.
- [10] Lianli Gao, Qilong Zhang, Jingkuan Song, Xianglong Liu, and Heng Tao Shen. Patch-wise attack for fooling deep neural network. In *European Conference on Computer Vision*, 2020.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.
- [13] Shixiang Gu and Luca Rigazio. Towards deep neural network architectures robust to adversarial examples. *arXiv preprint arXiv:1412.5068*, 2014.
- [14] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens Van Der Maaten. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*, 2017.

- [15] Jamie Hayes. On visible adversarial perturbations & digital watermarking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2018.
- [16] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *International Conference on Neural Information Processing Systems*, 2015.
- [17] Danny Karmon, Daniel Zoran, and Yoav Goldberg. Lavan: Localized and visible adversarial noise. In *International Conference on Machine Learning*, 2018.
- [18] Aishan Liu, Xianglong Liu, Jiaxin Fan, Yuqing Ma, Anlan Zhang, Huiyuan Xie, and Dacheng Tao. Perceptual-sensitive gan for generating adversarial patches. In *Proceedings of the AAAI conference on artificial intelligence*, 2019.
- [19] Aishan Liu, Jiakai Wang, Xianglong Liu, Bowen Cao, Chongzhi Zhang, and Hang Yu. Bias-based universal adversarial patch attack for automatic check-out. In *Proc. Eur. Conf. Comput. Vis.*, 2020.
- [20] Yan Luo, Xavier Boix, Gemma Roig, Tomaso Poggio, and Qi Zhao. Foveation-based mechanisms alleviate adversarial examples. *arXiv preprint arXiv:1511.06292*, 2015.
- [21] Chunchuan Lyu, Kaizhu Huang, and Hai-Ning Liang. A unified gradient regularization family for adversarial examples. In *2015 IEEE international conference on data mining*, 2015.
- [22] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Machine Learning*, 2018.
- [23] Muzammal Naseer, Salman Khan, and Fatih Porikli. Local gradients smoothing: Defense against localized adversarial attacks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019.
- [24] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie. Generative adversarial perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [25] Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, and R Venkatesh Babu. Nag: Network for adversary generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [26] A Senior, V Vanhoucke, P Nguyen, T Sainath, et al. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal processing magazine*, 2012.
- [27] Uri Shaham, Yutaro Yamada, and Sahand Negahban. Understanding adversarial training: Increasing local stability of supervised models through robust optimization. *Neurocomputing*, 2018.
- [28] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. 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*, 2016.

- [29] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [30] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014.
- [31] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.
- [32] Haotao Wang, Tianlong Chen, Shupeng Gui, Ting-Kuei Hu, Ji Liu, and Zhangyang Wang. Once-for-all adversarial training: In-situ tradeoff between robustness and accuracy for free. *Advances in neural information processing systems*, 2020.
- [33] Tong Wu, Liang Tong, and Yevgeniy Vorobeychik. Defending against physically realizable attacks on image classification. In *International Conference on Learning Representations*, 2020.
- [34] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. In *IJCAI*, 2018.
- [35] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. *Proceedings 2018 Network and Distributed System Security Symposium*, 2018.
- [36] Chenglin Yang, Adam Kortylewski, Cihang Xie, Yinzhi Cao, and Alan Yuille. Patchattack: A black-box texture-based attack with reinforcement learning. In *European Conference on Computer Vision*, pages 681–698. Springer, 2020.
- [37] Xiao Yang, Fangyun Wei, and Hongyang Zhang. Design and interpretation of universal adversarial patches in face detection. *ECCV*, 2019.
- [38] Xiulong Yang and Shihao Ji. Learning with multiplicative perturbations. In *International Conference on Pattern Recognition (ICPR)*, 2020.
- [39] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.