

PATCHGUARD++: EFFICIENT PROVABLE ATTACK DETECTION AGAINST ADVERSARIAL PATCHES

Chong Xiang, Prateek Mittal

Princeton University

cxiang@princeton.edu, pmittal@princeton.edu

ABSTRACT

An adversarial patch can arbitrarily manipulate image pixels within a restricted region to induce model misclassification. The threat of this localized attack has gained significant attention because the adversary can mount a physically-realizable attack by attaching patches to the victim object. Recent provably robust defenses generally follow the PatchGuard framework (Xiang et al., 2021) by using CNNs with small receptive fields and secure feature aggregation for robust model predictions. In this paper, we extend PatchGuard to PatchGuard++ for provably detecting the adversarial patch attack to boost both provable robust accuracy and clean accuracy. In PatchGuard++, we first use a CNN with small receptive fields for feature extraction so that the number of features corrupted by the adversarial patch is bounded. Next, we apply masks in the feature space and evaluate predictions on all possible masked feature maps. Finally, we extract a pattern from all masked predictions to catch the adversarial patch attack. We evaluate PatchGuard++ on ImageNette (a 10-class subset of ImageNet), ImageNet, and CIFAR-10 and demonstrate that PatchGuard++ significantly improves the provable robustness and clean performance.

1 INTRODUCTION

The adversarial patch attack against computer vision systems allows the adversary to arbitrarily manipulate pixels within a restricted region to induce model misclassification. Compared with the L_p norm bounded adversary, this localized threat model is more realistic in that the adversary can attach the patch to an object to mount a physical attack (Brown et al., 2017; Eykholt et al., 2018). To counter the threat of an adversarial patch, empirical defenses (Hayes, 2018; Naseer et al., 2019) were proposed first but later shown vulnerable to adaptive attackers. The fragility of empirical defenses has inspired provably robust defenses (Chiang et al., 2020; Zhang et al., 2020; Levine & Feizi, 2020; Xiang et al., 2021; Lin et al., 2021; Metzen & Yatsura, 2021) for making robust predictions.

Despite the significant progress in provably robust defenses, however, there is still a huge gap between the provable robust accuracy of defended models and the clean accuracy of undefended models. To narrow this gap, we relax the defense objective from making robust predictions to robustly detecting the attack (McCoyd et al., 2020). This objective is meaningful in applications where ML systems can alert users. For example, if an autonomous vehicle detects an adversarial patch attack, it can alert and return the control to the driver. Towards this end, we propose PatchGuard++ as an efficient and provably robust attack detection mechanism, which significantly improves both provable robust accuracy and clean accuracy compared with previous provably robust defenses.

In PatchGuard++, we first follow PatchGuard (Xiang et al., 2021) to use a CNN with small receptive fields for feature extraction. The small receptive field bounds the number of features that can be corrupted by an adversarial patch. Next, we apply a mask to all possible locations in the feature map and evaluate masked predictions. We empirically observe that the predictions of CNNs are generally invariant to partial feature masking, and thus all masked predictions of clean images usually reach a consensus. On the other hand, when an adversarial patch is attached, some masked predictions are affected by the adversarial perturbations while some are not (when all corrupted features are masked). This will lead to a disagreement in masked predictions. Our final step aims to catch this disagreement for provable attack detection. We evaluate our defense on ImageNette (fast.ai, 2020)

(a 10-class subset of ImageNet), ImageNet (Deng et al., 2009), and CIFAR-10 (Krizhevsky, 2009) datasets and demonstrate state-of-the-art defense performance for high-resolution images. Notably, we achieve a 95.3% clean accuracy and 92.2% provable robust accuracy on the 10-class ImageNette, against a 2%-pixel square patch (the patch can be anywhere on the image).

2 PROBLEM FORMULATION

Adversarial Patch Attack. We focus on the evasion attack against image classification models. Given an image $\mathbf{x} \in \mathcal{X} \subset [0, 1]^{W \times H \times 3}$, its true class label $y \in \mathcal{Y} = \{0, 1, \dots, N - 1\}$, and an image classification model $\mathcal{M}(\cdot)$, the goal of the attacker is to find an image $\mathbf{x}' \in \mathcal{A}(\mathbf{x})$ satisfying a constraint \mathcal{A} such that $\mathcal{M}(\mathbf{x}') \neq y$. The constraint $\mathcal{A}(\mathbf{x})$ is determined by the threat model. In this paper, we allow the adversary to arbitrarily manipulate pixels within one square contiguous region and place such restricted region *anywhere* on the image, which is similar to previous works (Levine & Feizi, 2020; McCoy et al., 2020; Xiang et al., 2021). Formally, we use a binary *pixel block* $\mathbf{p} \in \{0, 1\}^{W \times H}$ to represent the restricted region, where the pixels within the region are set to one. The constraint set $\mathcal{A}(\mathbf{x})$ can be expressed as $\{\mathbf{x}' = (\mathbf{1} - \mathbf{p}) \odot \mathbf{x} + \mathbf{p} \odot \mathbf{x}''\}$, where \odot refers to the element-wise product operator, and \mathbf{x}'' is the content of the adversarial patch.

Provably Robust Defense for Attack Detection. We consider the strongest adaptive attacker who has the perfect knowledge of the defense mechanisms as well as the white-box access to the model architecture and weights. We focus on the *provably robust* defense whose robust guarantee holds for *any* attacker considered in the threat model. The goal of *provable attack detection* is to make normal predictions on clean images but alert when an attack takes place.

3 PATCHGUARD++: PROVABLE ATTACK DETECTION

Defense Intuition. Our key observation is that clean images are generally prediction-invariant to partial feature masking while adversarial patched images are not. When benign features are masked, the model is likely to give the same prediction; when corrupted features are completely masked, the adversary cannot influence the model prediction, and the model prediction is likely to change. Therefore, we can move a mask, which is large enough to mask out all corrupted features, over the feature map, and take the inconsistent masked prediction as an attack indicator. In the remainder of this section, we will introduce our notation, the defense algorithm, and its provable analysis.

Notation. The defense takes an image \mathbf{x} as input and outputs prediction $\mathcal{M}(\mathbf{x})$ or issues an attack alert. The model $\mathcal{M}(\cdot)$ can be any CNN with small receptive fields (relatively small compared with the input image size). We use \mathcal{F} to denote the feature extractor of the model $\mathcal{M}(\cdot)$, and $\mathcal{F}(\mathbf{x})$ gives the feature map $\mathbf{u} \in \mathbb{R}^{W' \times H' \times C'}$. We will abuse the notation to let $\mathcal{M}(\mathbf{u})$ denote making prediction based on the feature map \mathbf{u} . We use $\mathbf{w} \in \mathcal{W} \subset \{0, 1\}^{W' \times H'}$ to denote a sliding window over the feature space, where elements within the window are set to ones and others are zeros.

Algorithm. We provide our defense pseudocode in Algorithm 1. We first get the prediction \bar{y}_0 and confidence \bar{c}_0 from the model with small receptive fields $\mathcal{M}(\mathbf{x})$, as well as feature map $\mathbf{u} = \mathcal{F}(\mathbf{x})$. Next, we try every window mask $\mathbf{w} \in \mathcal{W}$,¹ and get masked prediction \bar{y}_1 and confidence \bar{c}_1 from $\mathcal{M}((\mathbf{1} - \mathbf{w}) \odot \mathbf{u})$. If the confidence \bar{c}_1 is larger than confidence threshold τ and the prediction \bar{y}_1 is different from \bar{y}_0 , our defense detects an inconsistent prediction and returns `alert` (i.e., attack detected). If confidence \bar{c}_1 is smaller than τ , we consider this masked prediction “abstains” and will not affect our final prediction. Finally, if we find all possible non-abstained masked predictions are consistent, we assume there is no attack and return the original model prediction $\mathcal{M}(\mathbf{x})$.

Provable Analysis. We provide pseudocode for provable analysis in Algorithm 2. Given a clean image \mathbf{x} and its label y , we will evaluate all possible masked predictions. If all predictions do not abstain ($\bar{c}_1 > \tau$) and are correct ($\bar{y}_1 == y$), we will know at least one masked prediction is non-abstained and correct regardless of any attack strategy, including adaptive white-box attacks. Therefore, our defense (Algorithm 1) will either predict y correctly or alert for this certified image.

¹The size of window masks can be computed as $w = \lceil (p + r - 1) / s \rceil$, where p is an estimation of the upper bound of patch size, r is the receptive field size, and s is the stride size of receptive field (Xiang et al., 2021). This window mask size can ensure that at least one window masks out the all corrupted features.

Algorithm 1 PatchGuard++

Input: Image \mathbf{x} , feature extractor \mathcal{F} of model \mathcal{M} , confidence threshold τ , the set of sliding windows \mathcal{W}

Output: Prediction $\mathcal{M}(\mathbf{x})$ or alert

```

1: procedure PG2
2:    $\bar{y}_0, \bar{c}_0 \leftarrow \mathcal{M}(\mathbf{x})$   $\triangleright$  Predict w/o masks
3:    $\mathbf{u} \leftarrow \mathcal{F}(\mathbf{x})$   $\triangleright$  Extract features
4:   for each  $\mathbf{w} \in \mathcal{W}$  do
5:      $\bar{y}_1, \bar{c}_1 \leftarrow \mathcal{M}((\mathbf{1} - \mathbf{w}) \odot \mathbf{u})$   $\triangleright$  Mask
6:     if  $\bar{c}_1 > \tau$  and  $\bar{y}_1 \neq \bar{y}_0$  then
7:       return alert  $\triangleright$  Alert!
8:     end if
9:   end for
10:  return  $\mathcal{M}(\mathbf{x})$   $\triangleright$  No attack
11: end procedure

```

Algorithm 2 PatchGuard++ provable analysis

Input: Image \mathbf{x} , label y , feature extractor \mathcal{F} of model \mathcal{M} , confidence threshold τ , the set of sliding windows \mathcal{W}

Output: Whether \mathbf{x} has provable robustness

```

1: procedure PG2PROVABLEANALYSIS
2:    $\mathbf{u} \leftarrow \mathcal{F}(\mathbf{x})$   $\triangleright$  Extract features
3:   for each  $\mathbf{w} \in \mathcal{W}$  do
4:      $\bar{y}_1, \bar{c}_1 \leftarrow \mathcal{M}((\mathbf{1} - \mathbf{w}) \odot \mathbf{u})$   $\triangleright$  Mask
5:     if  $\bar{c}_1 < \tau$  or  $\bar{y}_1 \neq y$  then
6:       return False  $\triangleright$  Incor./Abst.
7:     end if
8:   end for
9:   return True  $\triangleright$  Provably robust!
10: end procedure

```

Remark. We note that Minority Report (MR) (McCoyd et al., 2020) also adopts a similar idea of masking and prediction consensus. However, MR uses pixel-space masks, which incurs an extremely large overhead (requires expensive model inference for each mask) and makes MR inapplicable to high-resolution images like ImageNet (Deng et al., 2009). In contrast, the feature-space masking in PatchGuard++ reduces the number of all possible masks and also reuses the expensive feature extraction computation, making PatchGuard++ easily scale to ImageNet images.

4 EVALUATION

In this section, we present the evaluation results of PatchGuard++ on different datasets and demonstrate state-of-the-art provable robust accuracy and clean accuracy on high-resolution images.

Datasets. We evaluate proposed defenses on three datasets: ImageNet (Deng et al., 2009), ImageNette (fast.ai, 2020), and CIFAR-10 (Krizhevsky, 2009). ImageNet is a 1000-class high-resolution (resized and cropped to 224×224) image classification benchmark; ImageNette is a 10-class subset of ImageNet. CIFAR-10 is a popular image benchmark dataset with a low resolution of 32×32 ; we rescale it to 192×192 as the input of PatchGuard++ as done in PatchGuard (Xiang et al., 2021).

Setup. We use BagNet-33 (Brendel & Bethge, 2019) with a small receptive field of 33×33 as the backbone CNN. Following previous works (Xiang et al., 2021; Metzen & Yatsura, 2021), we use 2%-pixel square patch for ImageNette and ImageNet, a 2.4%-pixel square patch for CIFAR-10.

Evaluation. In Table 1, we report the results for PatchGuard++ and also compare them with previous provably robust defenses. PG-Mask-BN (Xiang et al., 2021), PG-Mask-BN (Xiang et al., 2021), and BagCert (Metzen & Yatsura, 2021) are provably robust for making correct predictions while MR (McCoyd et al., 2020) and PatchGuard++ are provably robust for detecting an attack. For previous defenses, we report results that are available in its original paper. For PatchGuard++, we report results with different confidence threshold τ .

Observation: PatchGuard++ works extremely well on high-resolution images, significantly outperforming all previous defenses in terms of provable robust accuracy and clean accuracy. We report defense performance in Table 1. Since MR is computationally infeasible for high-resolution images, we only compare performance with robust prediction defenses on ImageNette and ImageNet. Notably, PatchGuard++ achieves a high clean accuracy of 96.1% and a high provable robust accuracy of 91.8% on ImageNette (setting $\tau = 0.6$), outperforming all previous defenses. On ImageNet, PatchGuard++ has a 6+% higher clean accuracy, and 13+% higher provable robust accuracy than all previous defenses when setting $\tau = 0.5$! Moreover, if we set $\tau = 0.3$ (not listed in the table), we can achieve a similar clean accuracy (56.1%) as PG-Mask-BN (54.6%) but a much higher provable robust accuracy (45.4%; 19.4% higher!).

Observation: PatchGuard++ also achieves comparable or better provable robust accuracy on low-resolution images. The comparison between MR and PatchGuard++ on CIFAR-10 shows that MR

Table 1: Clean and provable robust accuracy for different defenses (PG-Mask-BN, PG-Mask-DS, and BagCert are for robust prediction; MR and PatchGuard++ are for robust detection).

Dataset	ImageNette		ImageNet		CIFAR-10	
Accuracy	clean	robust	clean	robust	clean	robust
PG-Mask-BN (Xiang et al., 2021)	95.0	86.7	54.6	26.0	83.9	47.3
PG-Mask-DS (Xiang et al., 2021)	92.1	79.7	43.6	15.7	84.6	57.7
BagCert (Metzen & Yatsura, 2021)	NA	NA	46.0	22.9	86.0	60.0
MR (McCoyd et al., 2020)	computationally infeasible				92.4	43.8
					90.6	62.1
					78.8	77.6
PatchGuard++ ($\tau = 0.8$)	96.9	87.7	62.9	28.0	84.8	68.9
PatchGuard++ ($\tau = 0.7$)	96.6	90.2	62.7	32.0	82.5	71.7
PatchGuard++ ($\tau = 0.6$)	96.1	91.8	62.1	35.5	80.2	74.3
PatchGuard++ ($\tau = 0.5$)	95.3	92.9	60.9	39.0	78.0	76.3

has a slightly better provable robustness in certain settings. However, the computation cost of MR is $100\times$ higher than ours (making it inapplicable to images in higher resolutions). Moreover, we can see PatchGuard++ achieves a huge improvement in provable robust accuracy compared with previous robust prediction defenses.

Observation: Threshold τ balances the trade-off between clean accuracy and provable robust accuracy of PatchGuard++. As shown in Table 1, as we decrease τ , the clean accuracy drops while the provable robust accuracy increases (both due to more non-abstained predictions). We note that the provable robust accuracy increases much more quickly than clean accuracy drops, which also demonstrates the strength of our defense design.

5 RELATED WORK

Attacks. Brown et al. first proposed a universal targeted adversarial patch attack in the physical world to hijack the model prediction (Brown et al., 2017). Karmon et al. later proposed a more powerful LaVAN attack in the digital domain (Karmon et al., 2018). More following up works also showed that the localized adversary can achieve a robust physical attack against traffic sign recognition, human detection (Eykholt et al., 2018; Chen et al., 2018). Effective defenses against these realistic attacks are needed.

Defenses. Heuristic-based Digital Watermarking (Hayes, 2018) and Local Gradient Smoothing (Naseer et al., 2019) defenses were proposed to tackle the threat of localized attacks. However, they had been shown vulnerable to an adaptive attacker (Chiang et al., 2020). To provide a stronger security guarantee, Chiang et al. proposed the first provable defense against the localized patch attack using Interval Bound Propagation (IBP) (Gowal et al., 2019; Mirman et al., 2018). Afterwards, Clipped BagNet (CBN) (Zhang et al., 2020), De-randomized Smoothing (DS) (Levine & Feizi, 2020), Minority Report (MR) (McCoyd et al., 2020), PatchGuard (Xiang et al., 2021), BagCert (Metzen & Yatsura, 2021) were proposed to improve provable security guarantee against the adversarial patch attack. PatchGuard++ proposed in this paper can be viewed as a hybrid of PatchGuard (small receptive fields + feature masking) and MR (checking consensus of masked predictions), and it significantly outperforms previous provably robust defenses.

6 CONCLUSION

In this paper, we propose PatchGuard++ for provably detecting localized adversarial patch attacks. PatchGuard++ is an extension of PatchGuard, and achieves provable robustness by detecting inconsistency in all possible feature-masked predictions. Our evaluation on three benchmark datasets demonstrates that our defense significantly boosts clean performance and provable robustness on high-resolution images.

REFERENCES

- Wieland Brendel and Matthias Bethge. Approximating CNNs with bag-of-local-features models works surprisingly well on ImageNet. In *7th International Conference on Learning Representations (ICLR)*, 2019.
- Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. In *Conference on Neural Information Processing Systems Workshops (NeurIPS Workshops)*, 2017.
- Shang-Tse Chen, Cory Cornelius, Jason Martin, and Duen Horng (Polo) Chau. Shapeshifter: Robust physical adversarial attack on faster R-CNN object detector. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, volume 11051, pp. 52–68. Springer, 2018.
- Ping-Yeh Chiang, Renkun Ni, Ahmed Abdelkader, Chen Zhu, Christoph Studor, and Tom Goldstein. Certified defenses for adversarial patches. In *8th International Conference on Learning Representations (ICLR)*, 2020.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 248–255, 2009.
- Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1625–1634, 2018.
- fast.ai. ImageNette: A smaller subset of 10 easily classified classes from imagenet. <https://github.com/fastai/imagenette>, 2020.
- Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Arthur Mann, and Pushmeet Kohli. Scalable verified training for provably robust image classification. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4841–4850, 2019.
- Jamie Hayes. On visible adversarial perturbations & digital watermarking. In *2018 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops)*, pp. 1597–1604, 2018.
- Danny Karmon, Daniel Zoran, and Yoav Goldberg. LaVAN: Localized and visible adversarial noise. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pp. 2512–2520, 2018.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>, 2009.
- Alexander Levine and Soheil Feizi. (De)randomized smoothing for certifiable defense against patch attacks. In *Conference on Neural Information Processing Systems, (NeurIPS)*, 2020.
- Wan-Yi Lin, Fatemeh Sheikholeslami, jinghao shi, Leslie Rice, and J Zico Kolter. Certified robustness against physically-realizable patch attack via randomized cropping, 2021. URL <https://openreview.net/forum?id=vttv9ADGuWF>.
- Michael McCoyd, Won Park, Steven Chen, Neil Shah, Ryan Roggenkemper, Minjune Hwang, Jason Xinyu Liu, and David A. Wagner. Minority reports defense: Defending against adversarial patches. In *Applied Cryptography and Network Security Workshops (ACNS Workshops)*, volume 12418, pp. 564–582. Springer, 2020.
- Jan Hendrik Metzen and Maksym Yatsura. Efficient certified defenses against patch attacks on image classifiers. In *9th International Conference on Learning Representations (ICLR)*, 2021. URL <https://openreview.net/forum?id=hr-3PMvDpil>.
- Matthew Mirman, Timon Gehr, and Martin T. Vechev. Differentiable abstract interpretation for provably robust neural networks. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pp. 3575–3583, 2018.

Muzammal Naseer, Salman Khan, and Fatih Porikli. Local gradients smoothing: Defense against localized adversarial attacks. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1300–1307, 2019.

Chong Xiang, Arjun Nitin Bhagoji, Vikash Sehwal, and Prateek Mittal. Patchguard: A provably robust defense against adversarial patches via small receptive fields and masking. In *30th USENIX Security Symposium (USENIX Security)*, 2021.

Zhanyuan Zhang, Benson Yuan, Michael McCoyd, and David Wagner. Clipped bagnet: Defending against sticker attacks with clipped bag-of-features. In *3rd Deep Learning and Security Workshop (DLS)*, 2020.