



PatchCensor: Patch Robustness Certification for Transformers via Exhaustive Testing

YUHENG HUANG, University of Alberta, Canada

LEI MA, University of Alberta, Canada and The University of Tokyo, Japan

YUANCHUN LI, Institute for AI Industry Research (AIR), Tsinghua University, China

In the past few years, Transformer has been widely adopted in many domains and applications because of its impressive performance. Vision Transformer (ViT), a successful and well-known variant, attracts considerable attention from both industry and academia thanks to its record-breaking performance in various vision tasks. However, ViT is also highly nonlinear like other classical neural networks and could be easily fooled by both natural and adversarial perturbations. This limitation could pose a threat to the deployment of ViT in the real industrial environment, especially in safety-critical scenarios. How to improve the robustness of ViT is thus an urgent issue that needs to be addressed. Among all kinds of robustness, patch robustness is defined as giving a reliable output when a random patch in the input domain is perturbed. The perturbation could be natural corruption, such as part of the camera lens being blurred. It could also be a distribution shift, such as an object that does not exist in the training data suddenly appearing in the camera. And in the worst case, there could be a malicious adversarial patch attack that aims to fool the prediction of a machine learning model by arbitrarily modifying pixels within a restricted region of an input image. This kind of attack is also called physical attack, as it is believed to be more real than digital attack. Although there has been some work on patch robustness improvement of Convolutional Neural Network, related studies on its counterpart ViT are still at an early stage as ViT is usually much more complex with far more parameters. It is harder to assess and improve its robustness, not to mention to provide a provable guarantee.

In this work, we propose PatchCensor, aiming to certify the patch robustness of ViT by applying exhaustive testing. We try to provide a provable guarantee by considering the worst patch attack scenarios. Unlike empirical defenses against adversarial patches that may be adaptively breached, certified robust approaches can provide a certified accuracy against arbitrary attacks under certain conditions. However, existing robustness certifications are mostly based on robust training, which often requires substantial training efforts and the sacrifice of model performance on normal samples. To bridge the gap, PatchCensor seeks to improve the robustness of the whole system by detecting abnormal inputs instead of training a robust model and asking it

154

This work is done partially while Yuheng Huang was an intern with Yuanchun Li at Institute for AI Industry Research (AIR), Tsinghua University, China.

Yuanchun Li is the corresponding author and is also affiliated with Shanghai AI Laboratory, China.

Yuanchun Li is supported in part by the National Key R&D Program of China (Grant No. 2022YFF0604501) and NSFC (Grant No. 62272261). Yuheng Huang and Lei Ma are supported in part by Canada CIFAR AI Chairs Program, the Natural Sciences and Engineering Research Council of Canada (NSERC Grants No. RGPIN-2021-02549, No. RGPAS-2021-00034, and No. DGEER-2021-00019). Lei Ma's research is also supported by JST-Mirai Program Grant No. JPMJMI20B8 and JSPS KAKENHI Grants No. JP20H04168 and No. JP21H04877.

Authors' addresses: Y. Huang, University of Alberta, Edmonton, Alberta, Canada; email: yuheng18@ualberta.ca; L. Ma, University of Alberta, Edmonton, Alberta, Canada and The University of Tokyo, Tokyo, Japan; email: ma.lei@acm.org; Y. Li, Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China; email: liyuanchun@air.tsinghua.edu.cn. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1049-331X/2023/09-ART154 \$15.00

<https://doi.org/10.1145/3591870>

to give reliable results for every input, which may inevitably compromise accuracy. Specifically, each input is tested by voting over multiple inferences with different mutated attention masks, where at least one inference is guaranteed to exclude the abnormal patch. This can be seen as complete-coverage testing, which could provide a statistical guarantee on inference at the test time. Our comprehensive evaluation demonstrates that PatchCensor is able to achieve high certified accuracy (e.g., 67.1% on ImageNet for 2%-pixel adversarial patches), significantly outperforming state-of-the-art techniques while achieving similar clean accuracy (81.8% on ImageNet). The clean accuracy is the same as vanilla ViT models. Meanwhile, our technique also supports flexible configurations to handle different adversarial patch sizes by simply changing the masking strategy.

CCS Concepts: • **Software and its engineering** → **Software reliability**; • **Computing methodologies** → **Artificial intelligence**;

Additional Key Words and Phrases: Adversarial patch, neural networks, robustness certification, vision transformer, certified accuracy, deep learning

ACM Reference format:

Yuheng Huang, Lei Ma, and Yuanchun Li. 2023. PatchCensor: Patch Robustness Certification for Transformers via Exhaustive Testing. *ACM Trans. Softw. Eng. Methodol.* 32, 6, Article 154 (September 2023), 34 pages. <https://doi.org/10.1145/3591870>

1 INTRODUCTION

As the core to the success of many intelligent systems, **Deep neural networks (DNNs)** are widely adopted across various industrial applications and domains. With the substantial advance in **Deep Learning (DL)** techniques, there is also an increasing trend to deploy DL models in the physical world in safety-critical scenarios, e.g., face authentication on smartphones [74], driving assistance on autonomous cars [85], and intruder detection on surveillance cameras [15]. However, DNNs are well-known for their nonlinear property, which makes them highly sensitive to perturbation in their input space. Both natural perturbations [32] and evasion attacks [61, 76] could pose threats to these DL models. As an alternative to **Convolutional Neural Network (CNN)** in vision tasks and recurrent language model (e.g., Recurrent Neural Network and Long Short-term Memory) in language processing tasks, Transformers [18, 80], the current **state-of-the-art (SOTA)** DNNs, still suffer from this sensitivity. This limitation is one of the major threats to deploying DNNs in safety-critical scenarios.

To address this issue, both the **Software Engineering (SE)** community [19, 37, 44, 52, 72, 81, 82, 95, 96, 96, 97, 101, 102] and the **Artificial Intelligence (AI)** community [12–14, 30, 43, 54, 57, 59, 64, 66, 67, 88, 89, 99] are trying to improve the DL models' robustness through training, testing and verification. Related work can be roughly categorized into two types: (1) Improve model robustness in the digital domain. (2) Improve model robustness in the physical domain. The issue of digital domain robustness was first discovered by Goodfellow et al., and it is formalized as the well-known adversarial attack [26]. The model is defined as robust when its prediction is stable with a adversarially generated *global* L_p -bounded perturbation ($p = 1, 2, \infty, \dots$).

More recently, researchers have found that the attacks applied to the whole image with small perturbations are hard to implement in real world. The reasons are fourfold: (1) In practical scenarios, noisy physical environments can destroy the perturbations created by digital attack [48]; (2) the small magnitude perturbation may be missed by the camera because of sensor imperfections [20]; (3) it is challenging to create an attack modifying the whole background in the real world [20]; (4) printing small magnitude perturbations is less effective [20, 102]. Because of these practical concerns, patch attack (or physical attack) [8, 20, 70, 77, 102] is proposed, which is believed to be a

real threat and has a more realistic attacking scenario than digital attacks. Its perturbation is limited in a given region with no magnitude restriction. Along the line, researchers found that an attacker can fool the DNN in a self-driving car by attaching a printed sticker onto the road signs [21] or bypass face authentication by presenting a customized object in front of the camera [38]. In general, patch attack can be very helpful for the developer to understand the vulnerability of the model's robustness under the patch-like attacks/noises, so that to further propose techniques to counteract, since they can be highly relevant and linked to the safety and security issues. However, as the patch attack has a totally different attack model from the digital attack, the defense for the digital attack is usually hard to apply. With this real threat, it is thus urgent and necessary to propose a defense mechanism that is specifically designed for the patch attack.

With the advance of related study of the model's robustness against adversarial distribution shift, there is also an ongoing trade in studying the natural distributional robustness of DNNs. Its motivation is that the worst-case adversarial perturbations are unlikely to happen in reality, while the corruptions usually seen in the natural environment can seriously influence models' performance. Similar to the adversarial setting, natural corruption can also be applied to the whole image or a limited patch. Examples of the former can be gaussian noise or motion blur of cameras [32] in the entire image. For the latter, it can be noise restricted to a patch [29] or occlusion in object detection [98] and anomaly objects in the open-set video setting [3].

Regarding digital-domain robustness, the adversarial perturbations and natural perturbations usually cannot be analyzed in the same framework as the attack magnitude is restricted by L_p -bounded while the natural perturbations are generally unbounded. A recent study also points out that there is a tradeoff between adversarial and natural robustness [58]. However, the patch robustness of both adversarial and natural perturbations can actually be analyzed together. Both set no restriction for perturbation magnitude, and only the patch size is limited. In this sense, given the same patch region, if the robustness of the worst-case adversarial attack is guaranteed, then the robustness of natural perturbations is also ensured.

Although patch robustness is important as it is closely related to both adversarial and natural threats, related studies are still at the early stage. Various heuristic defenses [59, 64] have been designed to improve patch robustness. However, these heuristic approaches are unable to guarantee robustness against unseen adaptive attacks [13]. Some provable defenses were recently proposed for adversarial patches. Most of them [43, 57, 88, 99] aim to design a model that can guarantee consistent predictions on images with or without an adversarial patch. These approaches are mostly based on the insight that the prediction can be made by aggregating local features extracted from small independent receptive fields. The certified robustness is achieved by ensuring that the aggregated result will not be dominated by a small number of potentially adversarial receptive fields. As a result, they have to significantly compromise the prediction accuracy due to the lack of global features. One recent work [57] in this type can only achieve a certified accuracy of 26.0% on ImageNet under 2%-pixel adversarial patches.

Instead of attempting to hazardedly make imprecise predictions in the presence of abnormal patches, we argue that it might be more reasonable to warn the user if the input is potentially harmful. For example, imagine a DNN is used for driving assistance, it would be desirable for the model to show how reliable and precise it is when it gives a driving suggestion while informing the driver about the risk if there is a situation that the model can not handle. Certified detection is designed for this purpose. The Minority Reports Defense [54], PatchGuard++ [89], and ScaleCert [30] all seek to detect the presence of an adversarial patch by partially occluding the image around each candidate patch location and analyzing the predictions of all occluded images. However, the Minority Reports Defense cannot scale to complex high-resolution images (e.g., ImageNet) due to the training difficulty and heavy computation overhead, and PatchGuard++ and ScaleCert are based

on CNN backbones with small receptive fields (e.g., BagNet), which cannot sufficiently utilize the global feature. As a result, how to obtain accurate and scalable robustness certification against adversarial patches remains an open problem.

To this end, we introduce PatchCensor, a testing-based method to certify the robustness against adversarial patches that can flexibly and accurately scale to complex images and different adversarial patch sizes. PatchCensor falls in the category of certified detection and aims to warn the user whenever an abnormal patch appears. As a testing-based method PatchCensor is zero-shot (aka. requires no additional training) by directly reusing a pretrained ViT backbone for certification. By providing the provable guarantee for the worst-case patch attacks, our method could also be applied to the detection of the corrupted natural patch.

PatchCensor is explicitly designed for Transformer architecture because:

- Transformers have achieved state-of-the-art performance across both **Natural Language Process (NLP)** and **Computer Vision (CV)** tasks.
- Recent studies show that Transformers are more vulnerable against physical-domain adversarial patches [28], addressing the need to defend it.
- Prior work has shown that Transformer-based neural network architecture is robust against absent blocks [60, 62], making it more reasonable to design mutation rules that explicitly mask some blocks.

Although our mutation strategy is designed for Transformers and works at the abstract feature block level, we will focus on **Vision Transformer (ViT)** in the rest of the article, because (1) most related work is designed and evaluated on vision tasks, and it would be more reasonable to compare with them in the same domain, and (2) both adversarial patch attacks and corresponding defense are formally defined in vision tasks but not in the natural language domain.

In a ViT model, the input image is partitioned into small patches. The patches are fed into several Transformer encoder layers to exchange local information by attending to each other. The global feature is obtained after the Transformer layers and used to classify the whole image. Such an attention mechanism enables a convenient way to exclude local patches from an inference pass—a patch can be excluded by masking the attention of other patches toward it. The remaining patches can still interact with each other to obtain sufficient global information for prediction.

Given an input image of a machine learning model, PatchCensor checks whether the prediction of the model may be controlled by an adversarial patch by testing it with different mutations. In each mutation, we exclude a region of the image by manipulating the attention mask of ViT and check whether the mutated inference can produce the same result as the original prediction. The exhaustive testing is done if all possible locations of the adversarial patch are tested, and an input is certifiably robust if all the tests reach a consensus on the same prediction. By designing the optimal strategy of attention mask generation, PatchCensor can efficiently achieve full test coverage with a minimal number of mutations. Given a data distribution, the robustness certification of PatchCensor is obtained by running exhaustive testing on the samples in the distribution. Similar to most certification approaches [42, 87], the certified accuracy is calculated as the ratio of samples that are correctly predicted and certifiably robust.

We evaluate PatchCensor on popular datasets, including ImageNet and CIFAR-10, with adversarial patch sizes ranging from 0.4% to 25%. The results demonstrate that PatchCensor is able to achieve significantly higher clean accuracy and certified accuracy than existing certified defenses. In particular, PatchCensor is able to stably achieve a high test accuracy in the verified image subset under different adversarial patch sizes, although the ratio of verifiable inputs decreases under larger adversarial patches.

Finally, we discuss the advantages and disadvantages of PatchCensor over the traditional certified recovery approaches. We empirically demonstrate the superior ability of PatchCensor in dealing with images with a small **region of interest (ROI)**. We also highlight the difference between the statistical certificate obtained by PatchCensor and traditional certificates, which has rarely been discussed in literature before.

In summary, this article makes the following contributions:

- We originally propose a zero-shot, testing-driven certified defense against adversarial patches by utilizing the characteristics of Transformer-based DNN architecture.
- We implement and test our method on popular datasets, including ImageNet and CIFAR-10, and we demonstrate that our method can achieve state-of-the-art certified accuracy on the datasets.
- We discuss the fundamental differences between certified recovery and certified detection through experiments, which explains the better practicality of our approach.

To the best of our knowledge, PatchCensor is a very early work that investigates the machine learning software security certification from the combined view of Transformer and exhaustive testing. We demonstrate that this could be a promising direction with the advantage and practical usefulness compared with other state-of-the-art techniques, which potentially inspires future research studies to safeguard the security of data-driven intelligent software.

2 BACKGROUND AND RELATED WORK

In this section, we first briefly introduce the background of the ViT and related work about its robustness. Then, we discuss previous work that is relevant to safeguarding DNNs in two directions: (1) DNN testing, which aims to improve the safety, quality, and robustness of DL models through empirical testing; (2) Patch Attack defense, which aims to defend DL models against the worst-case adversarial threat.

2.1 Vision Transformer and Its Robustness

The ViT architecture [18] is proposed recently as an alternative to CNNs [75] in computer vision tasks. ViT is inspired by the scaling success of self-attention-based architectures (in particular Transformer [80]) on NLP tasks. In a ViT, an image is split into patches and converted to a sequence of linear embeddings of these patches, which is then fed into a Transformer as the input. Image patches are treated the same way as tokens (words) in an NLP application. By pre-training the model on large amounts of data, it is able to attain excellent results as compared to convolutional networks. Later studies [47, 93, 100] have shown that the Transformer architecture can achieve state-of-the-art performance on most CV tasks.

Despite the great success of ViT, its robustness against adversarial attacks is still in doubt. Researchers have found that ViT is robust toward occlusion and natural perturbations [5, 60, 62], while it is vulnerable to adversarial attacks [4, 22, 28, 53]. Bai et al. [4] claim that transformers are at least no more robust than CNNs when defending digital-domain adversarial perturbation, while Gu et al. [28] point out that ViT is significantly more vulnerable against physical-domain adversarial patches. Furthermore, Fu et al. [22] develop an attack framework called Patch-Fool that targets the self-attention mechanism of the ViT specifically. They also find that ViTs become weaker learners compared to CNNs under this setting, addressing the need for improving the robustness of ViT. In this work, we aim to design a certified adversarial patch defense based on ViT by utilizing its strength (robustness to occlusion) to mitigate its weakness (vulnerability against adversarial attacks).

2.2 DNN Testing

How to improve the quality and robustness of DNNs and ensure their correctness at the deployment stage remains to be an open problem. In recent years, we have witnessed ongoing efforts in the SE community trying to do quality assurance for DL models through testing. They can be roughly categorized into two types: (1) those that aim to assess the quality of DNNs before deployment by synthesizing test inputs and (2) those that aim to capture subtle behaviours exhibited in DL models at a single input level.

The former is motivated by the fact that testing data are expensive and often insufficient to assess the quality of DNNs. Motivated by this, some testing techniques are developed to synthesize test inputs to thoroughly test the DNNs [49, 51, 78, 84, 90, 94, 102]. The test generation process can be guided through a series of DNN coverage criteria [50, 63]. The most relevant one is Deep-Billboard [102], which aims to generate patch perturbations to test Autonomous Driving systems. However, these testing techniques are parallel to our work as they focus on evaluating the quality of DNNs at the development stage while we try to detect abnormal situations given a single input at the deployment stage.

Another line of work tries to probe the boundaries of DNNs' capability with a single test input. Abnormal behaviour analysis tries to analyze the models' behavior and understand when the model could yield erroneous prediction [37, 44, 52, 72, 81, 95, 96]. This analysis can be done by assessing the models' uncertainty [52, 95, 96], by detecting whether the input instance is **Out-of-Distribution (OOD)** or not [37, 44, 81], or by directly predicting the failure of the AI system [44]. Adversarial detection [19, 82, 97, 101] has a slightly different focus and tries to detect inputs that are adversarially tampered with. Our work is orthogonal to them as we aim to perform exhaustive testing for a given input to achieve a certified level guarantee, while most related work attempts to provide empirical results.

Our work is built on a specific technique called testing by mutation. It has been widely used in software engineering and has recently been adapted for testing complex DL-driven systems in both the development [33, 34, 46, 51, 65, 78, 90, 94] and deployment [82] stages. The key to this technique is to design suitable mutation operators carefully as different tasks (e.g., CV [65, 78, 90, 94] or NLP [33, 46]) and different targets (e.g., model input mutation [33, 46, 65, 78, 90, 94] or model mutation [34, 51]) may require totally different operators. PatchCensor apply exhaustive testing for mutated model inputs at the deployment stage. The mutation operator is designed at the abstract feature level. Although our discussion mostly centers around the context of computer vision, our proposed mutation technique is general and could also potentially be applied to other application scenarios with some adaptation.

2.3 Adversarial Patch Attack and Defense

Many studies have shown that DNNs could be vulnerable against adversarial attacks [26, 76]. Classical adversarial attacks are known as adversarial perturbations, aiming to generate a small perturbation added to the pixels of an image that can lead to misclassification.

Instead of the whole-image perturbations of classical adversarial attacks that may be unrealistic in the real world [13], adversarial patch attack performs pixel modification within a restricted region in the image, which is easily achievable in the physical environment (e.g., putting a sticker on the traffic sign, wearing a specially designed glass or T-shirt, etc.). Brown et al. [8] first demonstrated the feasibility of fooling an image classifier with a specifically crafted physical patch. Along this direction, numerous other approaches [21, 35, 39, 69, 102] have been proposed to achieve more practical and effective patch attacks, or test real-world AI-based software systems under the adversarial patches.




Fig. 1. An illustration of detection-based certified defense.

To better ensure the reliability of AI-based software, there also comes a recent trend to propose security-defending techniques against adversarial patches, including heuristic defenses and certified defenses. Heuristic defenses [12, 14, 59, 64, 66] are mostly based on the empirical understanding of existing attacks. Among them, empirical adversarial detection shares a similar setting with us. The problem is formulated as a binary classification problem. Some approaches utilize another machine learning model (such as DNN) to classify the inputs as natural or adversarial by directly analyzing the input [25, 27, 56] or utilizing features extracted from the neural networks [1, 10, 83]. Another line of work tries first to analyze the DNNs to understand their internals and leverage related knowledge for the detection [19, 45, 92]. Finally, some papers focus more on the input and try to detect adversarial inputs through input reconstruction [55, 68] or feature squeezing [91]. However, such empirical attacks may fail on strong adaptive attacks [9, 13] and are thus unreliable in safety-critical scenarios. Certified defenses, instead, aim to rigorously guarantee the robustness of a given model to patch attacks. It is worth noticing that different from whole-image perturbations, where adversarial and natural perturbations are hard to analyze in a uniform framework, it is possible to unify both kinds of perturbations in the patch setting. In other words, given a fixed patch size, once the robustness against the worst-case adversarial case is guaranteed, the safety for natural perturbations can then be obtained also. Due to the promising capability to defend strong adaptive adversaries, certified defenses have attracted many researchers' interest in recent years.

At a high level, existing certified defenses mostly fall into two categories: *certified recovery* [13, 43, 57, 67, 88, 99] and *certified detection* [30, 54, 89]. Certified recovery aims to recover the correct prediction on images with adversarial patches, while certified detection aims to detect adversaries with a provable detection rate, as shown in Figure 1. In the following, we briefly introduce three representative certification workflows that are closely related to ours, two for certified recovery and one for certified detection. Notice that the success of the patch attack comes from the non-linear property of the neural networks: a small local area can influence the prediction of the global input. As a result, these works share a similarity that they all try to isolate or block the influence of the adversarial patch through different perspectives.

(De)Randomized Smoothing (PatchSmoothing) [43] is one of the early attempts that belongs to the certified recovery category. Its main contribution is to select only part of the image as one input and perform multiple inferences across the whole image. There are two selection strategies in the article: block smoothing and band smoothing. The block smoothing chooses a square with size $s \times s$ as one input and enumerates all $h \times w$ positions where h is the height of the image, and w is the width of the image. When the center of the block hits the boundary of the image, PatchSmoothing wraps around (e.g., use the pixels in the other three corners if the center is positioned in one corner). After counting the inference results of all these squares, PatchSmoothing computes the difference between top-2 classes, and if

$$\Delta \geq 2(m + s - 1)^2, \quad (1)$$

then PatchSmoothing certifies the input with the label of top-1 class. Here Δ means the difference, m is the size of the adversarial patch, and s is the square size. This threshold is decided by the maximum number of patches the adversarial can control. Band smooth works similarly, but instead, choosing a column with size s as the input and the threshold is

$$\Delta \geq 2(m + s - 1). \quad (2)$$

PatchGuard [88] moves one step further. It utilizes a particular CNN called bagnet [7], which has small receptive fields. The restricted receptive fields limit the effects of the adversarial patch. Based on such a structure, PatchGuard proposes an element-wise linear aggregation method to replace the original unsecured fully-connected layer, which will mix both benign and adversarial patches' prediction results. The aggregation mechanism works as follows: It checks the prediction vectors of all the sliding windows with a predefined size (adversarial patch size). If a sliding window with the highest score with respect to one class exceeds a predefined threshold, then PatchGuard will mask that sliding window and uses the remaining window to predict one class. Such a mechanism is motivated by the observation that the adversarial patch will increase the score of a target class to an extremely high value to surpass the original prediction.

These two certified recovery methods both utilize the voter mechanism, while the former ensures that the adversary will not influence the majority, and the latter ensures the suspicious voter can't control the result. We detail these two methods that are parallel to our work, because, in this article, we want to discuss the difference between certified recovery and certified detection and provide both practical concerns and experiment results for our choice of certified detection.

Minority Report (MR) [54], however, is a representative work of certified detection. For every input, it also makes multiple inferences where each one is masked by an occlusion region $s \times s$ that is two pixels larger than a predefined adversarial patch. The occlusion region slides across the image (with size $h \times w$), which requires MR to collect results for $(h - s + 1)(w - s + 1)$ inference. MR only certifies the result if all predictions yield the same class. While in their implementation, this requirement is relaxed by first discarding the lowest score to tolerate outliers and providing certified results when the supported votes exceed a predefined threshold. Such relaxation will incur false negatives but increase the overall performance.

The focus of this article is certified detection. PatchCensor differentiates existing certified detection methods in three ways:

- PatchCensor is built on an exhaustive testing strategy and does not require extra effort to train a model with patch robustness.
- Certification of PatchCensor works on a range of patch sizes instead of one fixed size, which is more flexible.
- PatchCensor can be applied for patch sizes up to 25% (while most related work is evaluated on 2%-pixel attack), which could be more practical when detecting natural patch perturbations.

3 MOTIVATIONAL STUDY

Certified recovery and certified detection for adversarial patches both aim to improve the robustness of DNNs and provide a formal guarantee. In this section, we present preliminary experiments to show the shortcomings of certified recovery and our motivation behind choosing certified detection.

Most existing certification approaches for adversarial patch defense are focused on certified recovery. Such approaches are typically based on a special model (e.g., BagNet [7]) that gives predictions with small receptive fields. They mostly require splitting the input image into small regions, each of which is fed into the model to generate a local prediction. The defended prediction

is made by voting over all the local predictions, and the certification for an input sample is obtained if sufficient local predictions vote for the same label, ensuring that an arbitrary adversarial patch would not change the prediction. In other words, given \mathcal{Y} classes, the image is certified as $class_A$ if the following condition is satisfied:

$$\begin{aligned} &Vote_{class_A} - Vote_{class_B} > threshold, \\ &\text{where } Vote_{class_B} = \max_{i \in \mathcal{Y}, i \neq A} (Vote_{class_i}). \end{aligned} \quad (3)$$

The threshold differs between different methods.

Although certified recovery is usually expected to provide a more rigorous and formal guarantee, we believe that offering a lightweight and statistical guarantee for whether there is an attack may be more practical for real-world applications.

- First, one obvious shortcoming of certified recovery is its limited performance (e.g., only 26.0% certified accuracy on ImageNet for the most recent work [57]). This often makes such defense impracticable for real-world scenarios.
- Second, optimizing the model for both clean and adversarial inputs would inevitably decrease the model's accuracy in all existing certified recovery approaches, which may harm the user experience in normal settings. However, it is hard for the current certified recovery method to achieve moderate certified performance without such optimization. Besides, such optimization, usually done by re-training or even training from scratch, may take more time and computing resources.
- In addition, we observe that recovering the correct prediction under the presence of adversarial patches is difficult or even impossible in some cases, e.g., when the critical ROI is hidden by an adversarial patch. Such a mechanism may be difficult to deal with images with a small ROI, where most local voters may not get enough information to produce the correct prediction. With the advance of camera technology, it is more common to have high-resolution pictures nowadays, leading to a much smaller ROI in practice than the popular MNIST and CIFAR-10 datasets.

To further validate our observation and hypothesis, we performed a motivational study by conducting studies to investigate two state-of-the-art certified recovery approaches, i.e., **PatchGuard (PG)** [88] and **De-randomized Smoothing (DS)** [43].

Specifically, we rescale each image in CIFAR-10 to a smaller size and pad the image to the original size, as shown in Figure 3. In this way, we can control the ROI size of each image. Since the images after resizing are different from the original ones, we have to retrain the models. We train the **De-randomized Smoothing ResNet (DS-ResNet)** from scratch for 200 epochs, using the common settings in De-randomized Smoothing [43] and PatchGuard [88]. We evaluate these two methods with 2.4% adversarial patch size and record the distribution shift of the voting differences ($Vote_{Rank1_class} - Vote_{Rank2_class}$), as shown in Figure 4.

It is obvious that when the ROI becomes smaller, the margin between Rank-1 voting and Rank-2 voting becomes smaller. If the voting margin of a sample is below a threshold, then the model does not yield a certified result for the sample. The threshold can be understood more intuitively for randomized Smoothing (as PatchGuard is based on the weighted sum and the voting value is continuous), in which the threshold is computed by the Equation (2).

For example, if the m (target size to certify) is 5 (2.4% patch size) and the s (column size) is 4 (default setting), then the threshold is 16. In Figure 4, we can clearly see that the distribution of De-randomized Smoothing for size 20 has just shifted to the left of 16. As a result, there is a sharp drop in the certified accuracy at size 20. The detailed results can be found in Figure 9 (see Section 5).




Fig. 2. An illustration of the information used by different defense approaches for prediction and certification. Certified recovery approaches are based on small local regions. Our method is based on the whole image with a small region occluded.




Fig. 3. An illustration of image rescaling. The left one is the original image, and the right one is the image after rescaling.




Fig. 4. The voting distribution of PatchGuard and De-randomized Smoothing with different rescaling sizes. X-axis shows the value difference between top-2 voting. Y-axis is the number of samples that produce the voting result.

We additionally analyze test instances that are predicted correctly but unable to certify in the CIFAR-10 dataset to show that taking only local features as input may cause confusion for certified recovery methods. In other words, the difference between the top-2 classes of these instances does not exceed the threshold defined in these algorithms. The confusion matrix for those instances is shown in Figure 5, where the y-axis indicates the true class, and the x-axis indicates the ranked-2 class. We can clearly find that for both methods, dog versus cat, car versus truck, and frog versus bird are classes that are hard to certify. And these classes are sometimes hard to distinguish even for humans, given only a small patch of the image, as shown in Figure 2

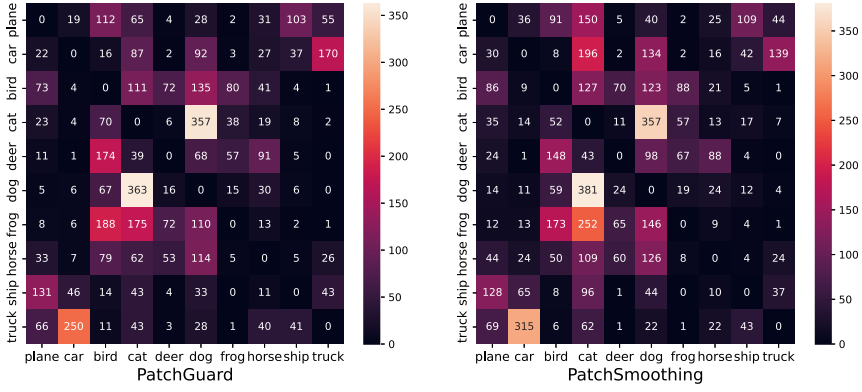


Fig. 5. Confusion matrix for test instances that are unable to certify but predicted correctly. Y-axis indicates the ranked-1 class and X-axis indicates the ranked-2 class.

The results of our preliminary study confirm that the certified recovery methods relying on local features could have fundamental limitations in dealing with small-ROI images, making them difficult to achieve reasonable certified accuracy in practice.

However, using as many patches as possible in the inference pass can provide sufficient information for the models to predict. This can improve clean accuracy, especially under a small ROI. However, increasing the number of patches in one inference will inevitably increase the chance of including malicious patches and, as a result, increase the difficulty for the defense. Our work aims to alleviate this dilemma by proposing a defense mechanism in a detection manner that can use more patches to improve both clean accuracy and certified accuracy.

4 PATCHCENSOR: CERTIFIED DETECTION VIA EXHAUSTIVE TESTING

In this section, we originally proposed a testing-based certified defense technique against adversarial patches, named PatchCensor. We first formulate the problem and formally define the properties that we intend to certify; then, we present our technique and make high-level theoretical analysis.

4.1 Problem Formulation

To be general, our work shares the same threat model as existing recent state-of-the-art certified defense work against adversarial patches [13, 43, 88, 99] (The attack and defense are both focused on image classification context.). We use $\mathcal{X} \subset \mathbb{R}^{W \times H \times C}$ to denote the distribution of images where each image $x \in \mathcal{X}$ has width W , height H , number of channels C . We take $\mathcal{Y} = \{0, 1, \dots, N-1\}$ as the label space, where the number of classes is N . We use $f: \mathcal{X} \rightarrow \mathcal{Y}$ to denote the model that takes an image $x \in \mathcal{X}$ as input and predicts the class label $y \in \mathcal{Y}$.

Attacker capability. The attacker can arbitrarily modify pixels within a restricted region, and this region can be anywhere on the image, even over the salient object. We assume that all manipulated pixels are within a square region, and the defender has a conservative estimate (i.e., upper bound) of the region size. Our technique can be generalized to other patch shapes, as long as the patch can be covered by a restricted rectangle, but we focus on square-shaped patches for simplicity.

Formally, we assume the attacker can arbitrarily modify an image x within a constraint set $\mathcal{A}(x)$. We use a binary pixel block $p \in \mathcal{P} \subset \{0, 1\}^{W \times H}$ to represent the restricted region, where the pixels within the region are set to 1. Then, the constraint set $\mathcal{A}(x)$ can be expressed as

$\{x' = (1 - p) \odot x + p \odot x'' \mid x, x' \in \mathcal{X}, x'' \in \mathbb{R}^{W \times H \times C}, p \in \mathcal{P}\}$, where \odot refers to the element-wise product operator, and x'' is the content of the adversarial patch.

Attack objective. We focus on adversarial patch attacks against image classification models. Given a deep learning model f , an image x , and its true class label y , the goal of the attacker is to find an image $x' \in \mathcal{A}(x) \subset \mathcal{X}$ such that $f(x') = y'$, where y' is an arbitrary incorrect class label defined by the attacker and $y' \neq y$.

Defense objective. The role of the defender is to design a defended model $g = (f, v) : \mathcal{X} \rightarrow \mathcal{Y} \times \{0, 1\}$, where $g(x) = (f(x), v(x))$, $f(x) \in \mathcal{Y}$ is the classification result, and $v(x) \in \{0, 1\}$ is the verification result indicating whether the prediction $f(x)$ can be verified (1 stands for “verified”). A verified inference means this inference is not influenced by any attacker, and we can trust the prediction made by the model.

The property to be certified. Based on the defense objective, we formulate the property to be certified as follows: Given a model M and a data distribution D , we certify that for any $x \in D$, the prediction $M(x)$ is correct and robust (cannot be falsified by an arbitrary attacker with ability A) with a probability of θ .

Certification for neural networks in the form of statistical guarantee is not a new idea. Most certified defending techniques in the digital attack domain [16, 42] also offer such a guarantee. However, this formulation may cause confusion, since most previous certification work (including certified detection, which is the topic and focus of this article) on patch attacks offers determined certification results. On the one hand, we find that for PatchCensor and all the previous certified detection work is, in reality, offering such statistical guarantee (discussed in Section 6), and we believe it is necessary and important to reformulate the problem. On the other hand, we argue that such a statistical guarantee is more practical in real-world applications. We will discuss more on this later.

Evaluation metric. The defender aims to improve the quality of the defended model by improving clean accuracy and certified accuracy. The clean accuracy acc_{clean} (the accuracy of f on the original dataset \mathcal{X} without considering verification) and the certified accuracy $acc_{\text{certified}}$ (the ratio of images in \mathcal{X} that are correctly and provably classified) are frequently used in prior work to describe the performance of certified defenses. Formally,

$$acc_{\text{clean}} = \mathbb{E}_{x \in \mathcal{X}}[l(f(x), y)], \quad (4)$$

$$acc_{\text{certified}} = \mathbb{E}_{x \in \mathcal{X}}[l(f(x), y) v(x)], \quad (5)$$

where l is the 0/1 error function.

In other words, the defender intends to maximize the ratio of verifiable inputs, measured by

$$r_{\text{robust}} = \frac{|\mathcal{X}_{\text{trust}}|}{|\mathcal{X}|} \quad (6)$$

and the model accuracy in the trust domain

$$acc_{\text{in-robust}} = \mathbb{E}_{x \in \mathcal{X}_{\text{trust}}}[l(f(x), y)], \quad (7)$$

so that it can retain high test accuracy and guarantee reliability on most inputs ($\mathcal{X}_{\text{robust}}$), while raising a warning if the input x is potentially malicious ($v(x) = 0$). The concept is similar to the selective prediction approaches that try to integrate a reject option in the neural network [23]. Such a rejection mechanism is helpful in many perception tasks in general where a fall-back solution is available. For example, a driving assistance model can ask the human driver to take over or perform a conservative operation when it's not confident about the current situation.

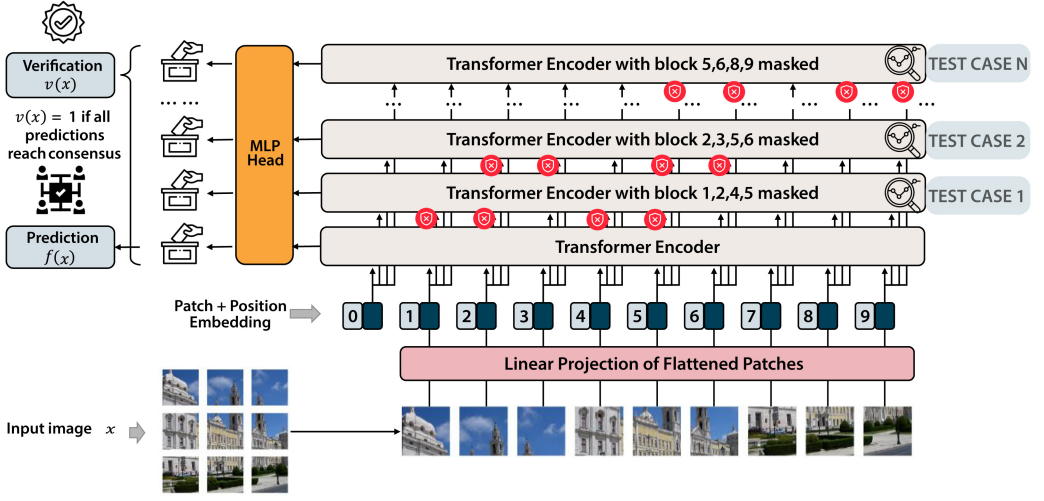


Fig. 6. The overall architecture and workflow of PatchCensor.

4.2 The Testing-based Defense

Our adversarial patch defense technique is designed upon the ViT architecture, by utilizing its input partition nature and self-attention mechanism.

The summarized workflow of PatchCensor are shown in Figure 6. It can be viewed as a paired function $g = (f, v)$, where f is a pretrained ViT model for prediction, and v is a verification function based on the ViT model. Given an input image x , PatchCensor produces a pair $(f(x), v(x))$, where $f(x)$ is the predicted class of x and $v(x)$ represents whether the prediction is verified.

In PatchCensor, the input image is split into non-overlapping patches as in vanilla ViT models. Suppose the input patch size is $P \times P$ in pixels, then the input image $x \in \mathcal{X} \subset \mathbb{R}^{W \times H \times C}$ is partitioned into a sequence of patches $\mathcal{P} = \{p_i | i = 1, 2, \dots, n\}$, where $n = n_w \times n_h = \frac{W}{P} \times \frac{H}{P}$ is the number of patches. Each patch p_i is then flattened, passed through a linear projection layer, and added a position embedding to generate a patch embedding q_i . A learnable [class] embedding q_0 is prepended to the sequence of embeddings, which is used to produce the class prediction after passing through later layers.

Specifically, the patch embeddings are fed into multiple parallel Transformer encoder layers $\mathcal{T} = \{t_0, t_1, \dots, t_k\}$ to exchange information between local patches. The encoders share the same weights of the ViT model f , while being used with different attention masks. Each Transformer encoder t_j produces an encoding of the [class] node, which is then passed through the MLP head to produce a class prediction y_j . We call the predictions produced with the masked attention maps (i.e., y_1, y_2, \dots, y_k) as *masked predictions*. To test whether there is an adversarial patch, we design a special mutation operator that directly masks the attention maps. The predictions with different masks can be seen as different mutations. The presence of the patch can be detected once we find anomalies among the test results. We call the predictions produced with the masked attention maps (i.e., y_1, y_2, \dots, y_k) as *mutations*. The class prediction $f(x)$ is produced by the Transformer encoder without mutation (i.e., y_0), which is equivalent to a direct inference using the base ViT model. The verification result $v(x)$ is produced by voting over all of the testing results (y_1 to y_k). We ensure that at least one of the mutations is benign (i.e., can completely mask the adversarial patch out), which vetoes the adversary's target output even if all other predictions are compromised. The prediction is verified if all Transformer encoders reach consensus (i.e., agreeing on the same class




Fig. 7. Number of ViT patches affected by different sizes of adversarial patches. P and W_{adv} represent the widths of the ViT input patch and adversarial patch, respectively.

prediction). As we enumerate all possible positions for the potential adversarial patch in this test generation process before masking, this test can exhaustively cover all corner cases, meaning that we complete a *full-coverage testing*.

Ideally, if testing can be done on all possible situations and these test cases are all passed (producing correct and consistent output), then the exhaustive testing would be equivalent to a rigorous verification. However, directly testing all possible adversarial patch positions is computationally infeasible and impractical considering its combinatorial complexity. Our mask strategy plays a role similar to the abstract set in abstract interpretation [17], providing a convenient (due to the natural robustness of self-attention architecture) and efficient (due to the reduced search space) way to achieve exhaustive full-coverage testing.

As a concrete example, we illustrate our method in Figure 6. Suppose we use a ViT model with 30×30 input resolution and 10×10 input patch resolution as the base model. An input image will be partitioned into 3×3 patches. To defend against adversarial patches with 5×5 resolution, we let each mask exclude 2×2 local patches, i.e., a square region with 20×20 resolution. By sliding the 2×2 mask over all local patches in the image, we can obtain four (2×2) possible mask locations and guarantee at least one of the mask locations can completely hide the 5×5 -pixel adversarial patch. A certified prediction is produced if the four masked predictions vote for the same class as the non-masked prediction.

4.3 Mutation Strategy

The certification capability of PatchCensor is achieved by testing with different mutations, while ensuring that at least one of the mutations can lead to an uninfluenced prediction (i.e., the corresponding attention mask can exclude the adversarial patch). The idea is analogous to Byzantine Fault Tolerance [11], where the benign voter can control the final result.

Each mutation leverages a different attention mask for inference. The key to designing the attention masks is to determine how many local input patches might be tainted by the adversarial patch (i.e., containing pixels belonging to the adversarial patch). Since the input partition plan is fixed and the adversarial patch is at an arbitrary location in the image, we need to consider the worst case, i.e., the maximum number of input patches that may be tainted by the adversarial patch.

Our masking strategy is based on the observation that, *in an image that is partitioned into non-overlapping $P \times P$ -pixel patches, an adversarial patch with $W_{adv} \times H_{adv}$ resolution can affect at most $N_W \times N_H$ input patches, where $N_W = \lceil \frac{W_{adv}}{P} \rceil + 1$ and $N_H = \lceil \frac{H_{adv}}{P} \rceil + 1$* . For example, as illustrated in Figure 7, if the width W_{adv} of a square adversarial patch is smaller than the input patch width P , then four input patches may be tainted (if the adversarial patch is at the joint of 2×2 input patches). Similarly, a square adversarial patch with $W_{adv} \in (2P, 3P)$ may affect at most 4×4 input patches.

We emphasize that our masking strategy here ensures that for any object at any location with any shape, as long as it can be restricted to a rectangle with shape $N_W \times N_H$, there must be a masking that can mask this object out. This flexibility makes it possible to detect both adversarial patches and abnormal objects in practice, while previous work often requires fine-tuning the model for a specific patch size.

Specifically, given the maximum rectangle shape $W_{adv} \times H_{adv}$ of the adversarial patch, we can compute the minimum size $N_W \times N_H$ for the attention mask. By sliding the mask over the whole input patch grid with a stride of 1, we can enumerate all k possible locations of the mask, where

$$k = \left(\frac{W}{P} - N_W + 1 \right) \times \left(\frac{H}{P} - N_H + 1 \right). \quad (8)$$

We can guarantee that at least one of the k masks can cover the arbitrary adversarial patch.

As we only mask a small proportion of the image and our base model is the powerful ViT, it is relatively easy for all voters to reach a consensus for a correct prediction on clean data.

4.4 Certification Analysis

In this subsection, we provide analysis to show that PatchCensor can achieve the defender's objective described in Section 3.

After obtaining the regular prediction of the ViT base model $f(x)$, the verification result of PatchCensor is obtained by

$$v(x) \triangleq \begin{cases} 1, & \text{if } f_1(x) = f_2(x) = \dots = f_k(x) = f(x), \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where $f_j(x)$ represents the prediction obtained by the ViT base model with the j th mask position on the attention map.

For any verified clean data point $x \in \mathcal{X}_{trust}$ and any adversarial example $x' \in \mathcal{A}(x)$, we need to ensure the adversarial patch is either ineffective or can be detected. Specifically, assuming x' can pass the verification, we have

$$f_1(x') = f_2(x') = \dots = f_k(x') = f(x'). \quad (10)$$

Based on our masking strategy (Section 4.3), at least one of the masked predictions will exclude the adversarial patch in x' , i.e.,

$$\exists j \in \{1, 2, \dots, k\}, \text{ s.t. } f_j(x') \neq f_j(x). \quad (11)$$

Since x is a verified input, we have $f(x) = f_j(x)$, and thus

$$f(x') \neq f_j(x') = f_j(x) = f(x), \quad (12)$$

meaning the adversarial patch attack on x' is ineffective.

However, it is still possible that the benign voter(s) make mistakes. This will eventually lead to an error as described in Section 6. However, such an error originates from the insufficient capability of the base model and is not influenced by the attacker. This error can also be estimated before model implementation as long as the data distribution for the benign patches (i.e., patches that do not contain OOD objects or are not adversarially influenced) is similar before and after implementation. Notice that the similar distribution assumption is one of the key assumptions in statistical machine learning theory [36, 79]

The estimated upper bound of the proportion of instances that the benign voters make mistakes is just $(1 - acc_{certified})$ on clean data. In other words, the probability of a correct and robust inference θ as we mentioned earlier is $acc_{certified}$. This estimated confidence can serve as an indicator of how much we can trust the model when the full-coverage testing is passed.

5 EVALUATION

To better understand the performance of PatchCensor and its relation with SOTA techniques from different angles, we design four research questions.

We first investigate the performance of our method and compare it with both certified recovery and certified detection methods in RQ1:

RQ1: How does the proposed system PatchCensor perform in terms of clean accuracy and certified accuracy?

This research question aims to offer an overview picture for the comparison of all the methods we have mentioned so far. To answer this research question, we select CIFAR-10 [40] and ImageNet [41] datasets under different adversarial patch sizes as our subject datasets for evaluation, which are also widely used in previous studies on adversarial patch defense. In addition to these two default evaluation datasets in related literature, we also conduct experiments on the **German Traffic Sign Recognition Benchmark (GTSRB)** dataset [71] and Food-101 dataset [6]. Resolution and number of instances usually reflect the difficulty of the dataset. We choose these four datasets to represent different levels of CV tasks. Specifically, CIFAR-10 contains 60,000 images with 32×32 resolution, GTSRB contains 51,839 images with 48×48 resolution, Food-101 contains 101,000 images with 384×384 resolution, and ImageNet contains 1,431,167 images with average resolution 469×387 . We report the clean accuracy and certified accuracy with 2.4% adversarial patch size on CIFAR-10 and 2% on ImageNet, GTSRB, and Food-101 dataset, following the setting in the previous work [13, 43, 54, 57].

After the general evaluation, we will verify the observations from our motivation study and discuss the difference between our methods (certified detection) and certified recovery when the area of ROI is small in RQ2:

RQ2: Does the proposed system PatchCensor achieve moderate certified performance even for images with a small area of the region of interest (ROI)?

Specifically, the aim of RQ2 is to discuss the reason behind choosing certified detection instead of certified recovery (as mentioned in the motivation study). Discussion on this decision is an important focus of this article, and RQ2 serves as a foundation for our further analysis in Section 6. It is also worth noting that we have compared the performance of PatchCensor with **advanced Minority Report (MR+)** in RQ1 and showed that PatchCensor achieves SOTA as a certified detection method. In RQ2, we aim to compare certified detection with certified recovery at the category level. Therefore, we select the most competitive ones from each category and did not include the second-best MR+ in this RQ.

After the discussion on certified detection and certified recovery, we then will offer an in-depth analysis of certified detection. Our evaluation is thus turned to the comparison between PatchCensor and the previous SOTA, MR+. One advantage of certified detection is its ability to defend the adversarially controlled patch with extremely large sizes (e.g., up to 25% of the whole image). So in RQ3, we want to investigate the performance of PatchCensor under such strong adversaries on CIFAR-10, GTSRB, Food-101, and ImageNet:

RQ3: How does the proposed PatchCensor perform under strong adversaries?

Specifically, we increase the adversarial patch size from 0.5% to up to 25%. Notice that it is possible to detect natural abnormal objects under such a setting. Also, when the size of dropped patches is large, fine-tuning may greatly influence the models' performance. So, we also discuss the influence of fine-tuning in this RQ. We do not report the results of certified recovery, because they are incapable of dealing with such large adversarial patches. In our experiments, the training is even hard to converge when the patch size is large.

Finally, performance improvement usually comes at a cost. We intend to understand what is the computation overhead for PatchCensor across different adversarial patch sizes compared with

previous state-of-the-art MR+. The overhead can be divided into two parts: the training time and the inference latency. In this RQ, we mainly discuss the inference latency as it is the most important factor in practical implementation:

RQ4: What is the computational overhead of PatchCensor compared with the state-of-the-art certified detection technique?

Notice that according to the mechanism we described in Section 4, the larger the target patch size, the fewer inference needed for PatchCensor. There is some trade-off in this relation as a large patch will yield worse certified accuracy but lower inference latency. To investigate such a relation, RQ4 includes experiments to guide choosing a suitable target patch size at runtime. We also provide experiments to demonstrate the influence of different ViT variants. Similar to RQ3, certified recovery methods will only provide trivial results when the patch size is large. Also, the fundamental certification goal is different, so we do not perform the evaluation in this RQ.

PatchCensor is implemented in Python based on a pre-trained ViT-Base model variant with 16×16 input patch size (ViT-B/16) [86], which achieves 81.8% test accuracy on ImageNet.¹

To support our large-scale evaluation, experiments were conducted on a Linux server with 4 Nvidia V100 GPUs. All the experiments take around 1,200 GPU hours. In the rest of this section, we summarize the key results for each of the studied research questions. Our code is released on the project website [2].

5.1 RQ1: Performance under Normal Setting

Experimental Settings. To evaluate the proposed PatchCensor under well-established settings, we compare PatchCensor (PC for short) with eight state-of-the-art certified adversarial patch defense techniques on CIFAR-10 and ImageNet as these two are, by default, the commonly-used datasets in related literature. The compared methods include four recovery-based methods and four detection-based methods. Recovery-based methods include **Interval Bound Propagation (IBP)** [13], **DS** [43], **PG** [88], and **BagCert (BC)** [57], and detection-based methods include **MR** [54], **PatchGuard++ (PG++)** [89], and **ScaleCert (SC)** [30]. Among them, MR is based on directly training (fine-tuning) a CNN model to classify the images with a square region occluded, thus can achieve higher clean and certified accuracy. However, the original design of MR needs to enumerate a lot of occlusion positions, which is computationally intensive for high-resolution images. Thus, we additionally implement a more advanced version of MR, named MR+, by porting our masking strategy to reduce its number of occlusion positions to the same as in Equation (8). We also replaced the CNN backbone of MR+ with a pre-trained ResNet50 [86] that has similar performance (81.1% accuracy on ImageNet) as the **ViT backbone (ViT-B/16)** used in PatchCensor.

After the experiments on CIFAR-10 and ImageNet, we further provide evaluations on GTSRB and Food-101. We include four methods: DS, PG, MR+, and PatchCensor. We choose these four methods because (1) they achieve competitive performance on CIFAR-10 and ImageNet, (2) they are representative methods in related areas, and (3) they are open-sourced.

Result. We first compare our approach with existing certified defense approaches in terms of clean accuracy and certified accuracy. As shown in Tables 1 and 2, our approach is able to obviously outperform the SOTA techniques with a clean accuracy of 81.8% and a certified accuracy of 67.1% on the challenging ImageNet with 32×32 -pixel adversarial patches. Due to the design of PatchCensor, its clean accuracy remains the same as the base ViT model, which can be even further improved by using better base models.

The results of detection-based approaches (including ours) are not directly comparable with the recovery-based approaches, because they are designed for different goals. However, we notice that

¹ All the source code, documents, and scripts are made available at the project website [2].

Table 1. The clean and Certified Accuracy of Different Certified Defenses on ImageNet and CIFAR-10

Method		ImageNet (2% patch size)		CIFAR-10 (2.4% patch size)	
		acc_{clean}	$acc_{certified}$	acc_{clean}	$acc_{certified}$
Recovery	Interval Bound Propagation (IBP) [13]	N/A		47.8	30.8
	De-randomized Smoothing (DS) [43]	44.4	14.0	83.9	56.2
	PatchGuard (PG) [88]	43.6	15.7	84.6	57.7
	BagCert (BC) [57]	45.2	22.9	86.0	60.0
Detection	Minority Reports (MR) [54]		N/A	78.8	77.6
				90.6	62.1
				92.4	43.8
	PatchGuard++ (PG++) [89]	62.9	28.0	84.8	68.9
		62.7	32.0	82.5	71.7
		62.1	35.5	80.2	74.3
		60.9	39.0	78.0	76.3
	ScaleCert (SC) [30]	58.5	55.4	78.9	75.3
	Minority Reports Adapted (MR+)	75.5	56.3	97.7	83.3
	PatchCensor (PC, our approach)	81.8	69.4	98.7	84.1

The numbers of IBP, DS, PG, BC, MR, PG++, and SC are directly copied from their paper. Note that the results of IBP and MR on ImageNet are not available, because they are computationally intensive or even infeasible on high-resolution images.

Table 2. The clean and Certified Accuracy of Different Certified Defenses on GTSRB and Food-101

Method		GTSRB (2% patch size)		Food-101 (2% patch size)	
		acc_{clean}	$acc_{certified}$	acc_{clean}	$acc_{certified}$
Recovery	De-randomized Smoothing (DS) [43]	52.73	15.97	50.15	17.33
	PatchGuard (PG) [88]	68.64	33.61	76.15	46.57
Detection	Minority Reports Adapted (MR+)	96.36	54.65	84.68	64.80
	PatchCensor (PC, our approach)	99.89	83.71	83.39	67.61

We fine-tune the models on GTSRB for 30 epochs and Food-101 for 60 epochs. All the models are fine-tuned with respect to the given patch size.

our approach was able to achieve a much higher certified accuracy than recovery-based methods, so it may be more practical to use in the real world. The main reason is that the certified detection is based on voting over predictions with a small region excluded, which can still provide sufficient global information, rather than in recovery-based approaches where each voter is based on a small local patch.

As compared to other SOTA certified detection approaches, PatchCensor could achieve higher certified accuracy on all four datasets. The results of clean accuracy are similar except for MR+ on Food-101, which means ViT-base performs slightly worse than ResNet-50. Notice that here we train both models with the same configuration (randomly occluded patches in the same size), to be relatively fair for comparison. The results indicate that the training configuration might not be optimal for ViT, which could still have some room for further improvement. Even though, under such a non-optimized configuration, the experiment results still confirm the effectiveness of our method in terms of certified accuracy. The result shows that the superior certified accuracy of PatchCensor comes not only from the better base performance of ViT but is also a result of




Fig. 8. Illustration of the annotation.

the combination of PatchCensor’s defense mechanism and ViT’s robustness against the absent patch.

Answer to RQ1: PatchCensor outperforms existing certified recovery methods by a large margin and also achieves higher clean accuracy and certified accuracy compared with other certified detection methods in most settings.

5.2 RQ2: Performance under Small ROI

Experimental Settings. In this research question, we want to evaluate the performance of PatchCensor under small ROI, as compared with the certified recovery methods mentioned in Section 3. The adversarial patch size we aim to certify in this RQ is 2.4%, which is 5×5 patch in a 32×32 image. We choose this patch size as it is the default setting in certified recovery work. We first perform rescaling for each image in the CIFAR-10 dataset so the ROI can be controlled. Then, we retrain the DS-ResNet on the rescaled image following the default setting in De-randomized Smoothing [43] and PatchGuard [88]. To further validate the PatchCensor can perform well on small ROI, we additionally design two experiments on the ImageNet visual object detection [73] and PartImageNet dataset [31]. The former includes bounding box annotations for the target object (as shown in Figure 8(a)) in the image and could serve as an indicator of how large the ROI is. We compute the ROI by dividing the bounding box’s area by the whole image’s area. The latter includes more fine-grained per-pixel part annotations, and each object is partitioned into smaller parts (as shown in Figure 8(b)). As current certified recovery methods usually rely on local features for inference and verification, the area of object parts could also stand for ROI. The ROI of PartImageNet is obtained by dividing the area of the largest part by the area of the whole image. The model architectures we used for all the methods here remain the same with RQ1.

Result. For the rescaled CIFAR-10 experiment, detailed result is shown at Table 3. We also plot the changing trend of $acc_{\text{certified}}$ and $acc_{\text{in-robust}}$ for PatchCensor and other two techniques in Figure 9.

One interesting observation is that certified recovery techniques (both the PatchGuard and De-randomized Smoothing) experience a quick drop for both r_{robust} and $acc_{\text{in-robust}}$ when the image size scales from 32×32 to 20×20 . The r_{robust} of De-randomized Smoothing even drops to nearly 0 (1.24%) with the rescaling size 20, unable to give certified prediction anymore. The reason behind this is the voting mechanism discussed in Section 3. While PatchCensor is able to remain a high certified accuracy even with very small ROI.

For experiments on ImageNet object detection and PartImageNet, we first sort the image according to ROI in ascending order and compute certified accuracy in 5-quantiles. The result for the former is shown in Figure 10, and the result for the latter is shown in Figure 11. It can be

Table 3. The clean Accuracy, Certified Accuracy, and Accuracy in Trust Domain Achieved by PatchCensor (PC), PatchGuard (PG), and De-randomized Smoothing (DS) for 2.4%-pixel Adversarial Patch Attack

Rescaling size	acc_{clean}			$acc_{certified}$			$acc_{in-robust}$		
	PC	PG	DS	PC	PG	DS	PC	PG	DS
32 (original size)	98.8	83.35	83.35	88.85	53.23	51.85	99.83	97.87	97.08
28	98.3	80.80	80.29	84.48	44.66	42.19	99.85	96.92	97.06
26	97.81	77.47	77.65	80.44	38.24	35.94	99.79	96.3	96.43
24	97.79	77.05	76.85	77.94	33.21	30.96	99.76	93.79	96.99
22	97.14	76.26	75.63	71.45	28.70	23.24	99.68	92.11	97.44
20	96.67	69.38	67.35	66.56	10.81	1.03	99.75	78.73	83.06

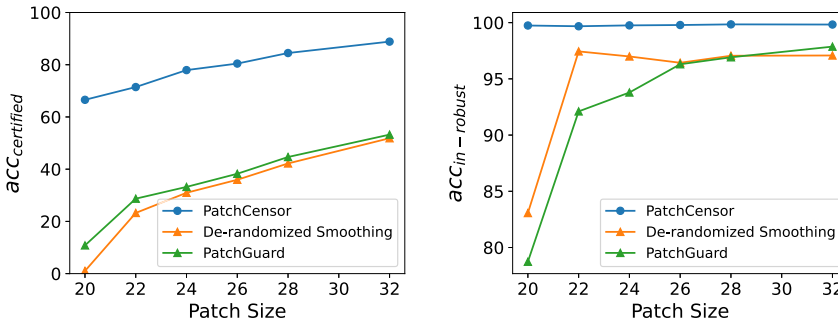


Fig. 9. Certified accuracy and accuracy in the robust domain of PatchCensor, PatchGuard, and De-randomized Smoothing on rescaled CIFAR-10.




Fig. 10. Certified accuracy and accuracy in the robust domain of PatchCensor, PatchGuard, and De-randomized Smoothing on ImageNet.

observed that PatchCensor can still have 51.0% certified accuracy on ImageNet object detection and 72.6% on PartImageNet in the 0-20% region, while PatchGuard and De-randomized Smoothing can only yield 18.7% (24.2%), and 11.0% (16.2%) certified accuracy, respectively.

It is worth noticing that $acc_{in-robust}$ of PatchCensor has some overlaps across different ROI on ImageNet object detection and PartImageNet. This may be due to the fact that pre-trained models on ImageNet may utilize other features in the image, such as texture instead of shape, as pointed out by Geirhos et al. [24]. In such cases, using only local features may be enough. However, this can be a kind of overfitting. Our controlled experiments on CIFAR-10 have no such issue. Despite




Fig. 11. Certified accuracy and accuracy in the robust domain of PatchCensor, PatchGuard, and De-randomized Smoothing on PartImageNet.

this, extra information such as texture could provide performance gain for certified recovery, and we will discuss more in Section 6.

Answer to RQ2: PatchCensor is able to achieve moderate performance when ROI is small. It can even maintain relatively good performance under small ROI at more complex and challenging ImageNet dataset.

5.3 RQ3: Performance under Strong Adversaries

Experimental Settings. In this research question, we evaluate the certification performance of PatchCensor and MR+ with large adversarial patches whose size is up to 25.0%. We ignore certified recovery work, because they only focus on small patches and are unable to handle large patches. On the one hand, the certification technique yields few certified instances once the patch size becomes large. On the other hand, it is much harder to train the model under such a setting. However, our defense and other SOTA certified detection work focus mainly on masking strategy. It is possible to certify instances even when the adversarial patches are large in size. Being able to certify inputs with large patch sizes would also make it possible to detect abnormal scenarios other than the adversarial setting, such as natural perturbation [29], occlusion in object detection [98], and open-set detection [3]. As natural objects are often larger than commonly evaluated 2.0% patch size. In the experiments, we select CIFAR-10, GTSRB, Food-101, and ImageNet as the subject datasets for evaluation. For the models, ViT-B/16 (the same architecture in RQ1, RQ2) and ResNet50 pre-trained on ImageNet are selected for PatchCensor and MR++ accordingly. Additionally, as we target a large adversarial patch size in this RQ, fine-tuning on randomly dropped patches can potentially improve the certified accuracy for both methods. As a result, we report the performance with/without such fine-tuning to provide a more comprehensive analysis of both methods. This leads to eight groups of experiments in total as shown below.

Result. We tested the defense effectiveness of PatchCensor and MR+ under different adversarial patch sizes on four datasets, with a total of 112 configurations. The results for for CIFAR-10 are summarized in Tables 4 and 5, results for GTSRB are summarized in Tables 6 and 7, results for Food-101 are summarized in Tables 8 and 9, and the results for ImageNet are summarized in Tables 10 and 11.

For all four datasets, PatchCensor achieves better performance on the most important certified accuracy metric than MR+ in most of the settings. The maximum difference is 40.73% on CIFAR-10 without fine-tuning, and the minimum difference is 2.57% on GTSRB without fine-tuning. The only exceptions are on the CIFAR-10 dataset with fine-tuning, where MR+ surpasses PatchCensor

Table 4. The Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on CIFAR-10 under Different Adversarial Patch Sizes **without Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	98.74	96.80	94.57	79.48	94.89	79.84	99.66	99.55
2.0%	98.74	96.81	90.13	68.52	90.29	68.78	99.82	99.62
4.6%	98.74	96.80	84.12	54.90	84.24	55.07	99.86	99.69
8.2%	98.74	96.79	76.20	41.37	76.28	41.53	99.90	99.61
12.8%	98.74	96.81	67.41	31.32	67.46	31.53	99.93	99.33
18.4%	98.74	96.81	58.90	23.79	58.97	24.01	99.88	99.08
25.0%	98.74	96.79	48.52	18.50	48.57	18.63	99.90	99.30

Table 5. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on CIFAR-10 under Different Adversarial Patch Sizes **with Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	98.77	97.61	95.70	90.59	96.05	90.97	99.64	99.58
2.0%	98.80	97.73	92.84	86.77	93.06	87.00	99.76	99.74
4.6%	98.81	97.70	88.29	83.30	88.42	83.46	99.85	99.81
8.2%	98.77	97.51	82.90	79.78	83.04	79.82	99.83	99.82
12.8%	98.76	97.35	75.52	74.91	75.99	75.04	99.91	99.83
18.4%	98.76	96.90	66.68	71.03	66.77	71.19	99.87	99.78
25.0%	98.75	95.57	55.89	65.60	55.95	65.77	99.89	99.74

Table 6. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on GTSRB under Different Adversarial Patch Sizes **without Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	97.43	92.49	70.26	36.13	71.01	36.32	98.94	99.48
2.0%	97.43	92.48	40.89	22.04	41.31	22.06	98.97	99.93
4.6%	97.43	92.48	26.04	11.94	26.08	11.94	99.85	100
8.2%	97.43	92.49	18.39	7.19	18.42	7.21	99.87	99.78
12.8%	97.43	92.48	12.49	5.65	12.50	5.69	99.87	99.30
18.4%	97.43	92.49	8.99	4.73	9.02	4.79	99.74	98.68
25.0%	97.43	92.49	6.71	4.14	6.73	4.22	99.76	98.12

when the patch size is larger or equal to 18.4%. From the result, we infer that CNN, specifically fine-tuned for patch drop, can be more robust than ViT on a simpler dataset such as CIFAR-10.

As for the clean accuracy, PatchCensor is higher than MR+ on nearly all settings (ranging from 1.18% on ImageNet to 8.31% on fine-tuned GTSRB) except for Food-101 with fine-tuning. Still, even under such a circumstance, PatchCensor has higher certified accuracy than MR+ across all adversarial patches on Food-101.

Interestingly, both methods have the worst performance on the GTSRB dataset when the adversarial patch size is large. It can be surprising that the certified accuracy of GTSRB is even lower than that of ImageNet, which is believed to be far more complex than GTSRB. Also, fine-tuning plays an important role in the performance, giving 26.87% improvement for PatchCensor and 19.02%

Table 7. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on **GTSRB** under Different Adversarial Patch Sizes **with Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	96.84	91.31	83.10	56.79	83.97	58.42	98.96	97.21
2.0%	97.09	92.53	70.33	48.35	71.21	49.43	98.77	97.81
4.6%	96.88	91.09	61.29	36.48	61.88	37.61	99.05	97.01
8.2%	97.65	91.23	49.29	27.78	49.89	28.35	98.79	97.96
12.8%	97.32	89.96	43.93	21.77	44.36	22.33	99.02	97.48
18.4%	96.52	90.15	35.61	18.38	36.08	18.65	98.71	98.51
25.0%	96.33	88.02	28.31	15.39	28.86	15.86	98.08	97.05

Table 8. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on **Food-101** under Different Adversarial Patch Sizes **without Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	86.46	84.45	77.08	66.28	81.04	68.62	95.11	96.59
2.0%	86.46	84.46	72.30	60.38	74.91	61.85	96.52	97.63
4.6%	86.46	84.46	66.60	53.36	68.26	54.38	97.56	98.12
8.2%	86.46	84.46	60.02	45.89	61.15	46.55	98.15	98.58
12.8%	86.46	84.46	52.82	38.71	53.61	39.17	98.52	98.82
18.4%	86.46	84.46	44.93	31.71	45.52	32.05	98.70	98.94
25.0%	86.46	84.46	36.61	24.74	37.09	25.03	98.73	98.83

Table 9. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on **Food-101** under Different Adversarial Patch Sizes **with Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	83.28	84.77	72.56	70.57	77.66	73.67	93.44	95.79
2.0%	83.39	84.68	67.61	64.80	70.93	66.86	95.32	96.91
4.6%	84.02	84.82	63.35	59.16	65.42	60.40	96.84	97.95
8.2%	84.80	84.47	58.65	53.24	59.96	54.18	97.81	98.27
12.8%	84.90	85.06	53.53	48.58	54.42	49.21	98.37	98.73
18.4%	84.97	85.30	47.22	44.33	47.92	44.83	98.55	98.88
25.0%	85.62	85.33	42.32	38.57	42.81	39.02	98.85	98.84

improvement for MR+ on certified accuracy. The reason may be that, compared with ImageNet, the target object in GTSRB occupies most of the image. It is more likely that the dropped patch is inside the target object and thus has more effect on the data distribution, interfering with the model's prediction. When the size of the dropped patch becomes large, important details may be missed, causing a sharp decrease in performance. As such, we suggest fine-tuning the models under a similar situation.

For other subject datasets we evaluated besides the GTSRB dataset, the fine-tuning process does not influence the performance of PatchCensor too much. The average increase after fine-tuning for PatchCensor in certified accuracy across all patches is 5.39%, -0.73%, and 4.44% for CIFAR-10, Food-101, and ImageNet, respectively. However, MR+ relies heavily on fine-tuning. The corresponding increase in certified accuracy is 33.23%, 19.02%, and 4.72% for CIFAR-10, Food-101, and ImageNet.

Table 10. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on ImageNet under Different Adversarial Patch Sizes **without Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	81.8	80.62	72.0	62.68	80.3	68.27	89.73	91.81
2.0%	81.8	80.62	67.2	56.02	73.5	60.24	91.35	92.99
4.6%	81.8	80.62	61.9	49.52	67.0	52.75	92.44	93.88
8.2%	81.8	80.62	56.4	42.94	60.4	45.46	93.29	94.46
12.8%	81.8	80.62	50.5	36.16	53.7	38.22	93.96	94.62
18.4%	81.8	80.62	44.1	30.05	46.7	31.70	94.57	94.78
25.0%	81.8	80.62	37.1	24.07	39.1	25.37	94.80	94.88

Table 11. Clean Accuracy and Certified Accuracy Achieved by PatchCensor and MR+ on ImageNet under Different Adversarial Patch Sizes **with Fine-tuning**

Max Adv Patch Size	acc_{clean}		$acc_{certified}$		r_{robust}		$acc_{in-robust}$	
	PC	MR+	PC	MR+	PC	MR+	PC	MR+
0.5%	82.70	75.59	73.67	61.84	81.88	70.99	89.98	87.12
2.0%	82.73	75.51	69.41	56.31	75.92	63.05	91.43	89.32
4.6%	82.67	75.49	64.97	51.23	70.27	56.56	92.45	90.57
8.2%	82.66	76.34	60.54	48.17	64.92	52.55	93.26	91.66
12.8%	82.55	76.20	55.38	43.39	59.08	46.99	93.73	92.35
18.4%	82.51	75.81	50.66	38.77	53.76	41.63	94.23	93.12
25.0%	82.49	75.50	45.62	34.79	48.24	37.34	94.56	93.19

Furthermore, the certified accuracy of PatchCensor without fine-tuning is higher than that of MR+ with fine-tuning on 16 of 21 cases on these three datasets. We can thus conclude that PatchCensor can yield moderate performance on datasets similar to these three, even without fine-tuning. This means that the pre-trained models can be directly combined with PatchCensor to increase its patch robustness.

Meanwhile, it is interesting to notice that the accuracy of PatchCensor in the trust domain ($acc_{in-trust}$) remained high (even slightly increases) under larger adversarial patch sizes for all datasets. This means that PatchCensor can retain high usability under strong adversaries—even though PatchCensor may raise more warnings (i.e., report more unverifiable input images) when defending against stronger adversaries, it can still promise a high accuracy when it gives verified predictions.

Answer to RQ3: PatchCensor outperforms MR+ in most scenarios on CIFAR-10, GT-SRB, Food-101, and ImageNet datasets. It also retains high accuracy in the robust domain ($acc_{in-robust}$) even with large adversarial patches. Being capable of detecting large abnormal patches makes it possible to use PatchCensor in more complex natural perturbation settings.

5.4 RQ4: Overhead of PatchCensor

Experimental Settings. In this research question, we investigate how much computation overhead PatchCensor may incur. Typically, there are two kinds of overhead for DNN models when conducting the certification. One is overhead for retraining/fine-tuning the model for the designed property. The other is the certification latency.

Table 12. Chosed ViT Variants for Latency Measurement The base model is what we use for evaluation in RQ1, RQ2, and RQ3

ViT Variants	Number of Parameters	Clean Accuracy	Patch Size	Input size
ViT_s16_224	22,050,664	80.5%	16	224
ViT_b16_224 (base)	86,567,656	80.9%	16	224
ViT_l16_224	304,326,632	82.2%	16	224
ViT_b32_224	88,224,232	80.8%	32	224
ViT_b32_384	88,297,192	80.8%	32	384




Fig. 12. Certified accuracy and latency for ViT variants of different model sizes.

For training overhead, we have already provided results and analysis with/without fine-tuning on four datasets in RQ3. For CIFAR-10, Food-101, and ImageNet, PatchCensor can yield moderate results without any fine-tuning efforts. Another particular feature of PatchCensor regarding the training overhead we want to emphasize here is that for both fine-tuning and certification, PatchCensor targets adversarial patches whose size is in a specific range (i.e., for all patches whose size is smaller than a threshold). In contrast, MR and most other patch certification methods only target adversarial patches with only one size.

For verification latency, we measured the latency of PatchCensor to verify an input image against different adversarial patch sizes. Here, we compare our technique with the previous SOTA MR+, as mentioned in the design goal of RQ4.

Results. We measure the verification latency of PatchCensor to see whether it is acceptable in this RQ. As ViT has many variants, and the latency for them may have a high variance, we independently provide results on each of them. We choose ViT_s16_224 (small), ViT_b16_224 (base), and ViT_l16_224 (large) for investigating the influence of model size. We choose ViT_b16_224 (base), ViT_b32_224 (base), and ViT_b32_384 (base) for investigating the influence of patch and input size. Model details are shown in Table 12. We additionally provide the results on clean accuracy and certified accuracy evaluated on 1,000 randomly sampled images from ImageNet without fine-tuning for defense. We hope this can serve as a guide for deployment in practice.

The results of the performance regarding the model sizes are shown in Figure 12. It is surprising that the smallest model achieves both the best certified accuracy and lowest latency (similar latency with MR+) among all the three variants. The reason may be that this smallest model does not overfit the dataset compared to the other two bigger models and is thus more robust toward masking. A similar result is also reported in Reference [60], where the smallest ViT has a high accuracy for random patch drop when the information loss is high. For the other two variants, stronger backbones produce slightly higher certified accuracy as we expected.

The results of the performance regarding the input and patch sizes are shown in Figure 13. These three variants have similar certified accuracy with quite different latency. Given the same input




Fig. 13. Certified accuracy and latency for ViT variants of different input and patch sizes.

size, a higher patch size for the base model will drastically lower the latency. The ViT_b32_224 achieves even lower latency than MR+. However, a smaller base patch will provide a more subtle control of the certified patch size. As PatchCensor certifies the patch *in range*, its certified accuracy is in the format of the step function, such as 0–16, 16–32, 32–48, ... for base patch equals to 16 and 0–32, 32–64, 64–96, ... for base patch equals to 32. If one wishes to certify a patch with a size 47, then it is better to choose a small base patch (32–48) rather than a big base patch (32–64), because the former can yield a better certified accuracy.

The results demonstrate that PatchCensor is better when using a smaller model with a large base patch. Under such a scenario, PatchCensor can achieve lower latency than MR+ while providing higher certified accuracy on a complex dataset.

Still, we want to emphasize that PatchCensor is not restricted to any kind of ViT variant. Even with a large ViT with a small base patch, the higher latency incurred by PatchCensor is also a meaningful tradeoff in some realistic scenarios because of the various benefits it provides, such as no need for fine-tuning and higher certified accuracy. PatchCensor is also a flexible detection framework that can work parallel with other methods. For real-time scenarios where the latency constraint is tight, PatchCensor can be combined and play as a nice complement to other low-latency algorithms, where it focuses on those critical frames that need rigorous analysis.

Answer to RQ4: PatchCensor requires much less effort in the training (fine-tuning) process and can yield good performance even without fine-tuning. PatchCensor can provide high certified accuracy with low latency when the base model is small, or its patch size is large. Even when defending a large model, we argue the higher latency is an acceptable compromise to the advantage it brings.

6 DISCUSSION

In this section, we further discuss the difference between certified detection and certified recovery. We first discuss a scenario in which the defense of certified detection may be “bypassed.” This discussion is ignored in previous work and could potentially misguide the readers. Still, we argue that this kind of special case will not influence the practicability of certified detection. Then, we discuss one possible direction for improving certified recovery.

6.1 A Special Case for Certified Detection

The certified property in certified recovery and certified detection is different. The former follows the convention defined in digital adversarial defense literature: a model is said to be robust if its inference remains the same in the neighbourhood of one input. The only difference is that the

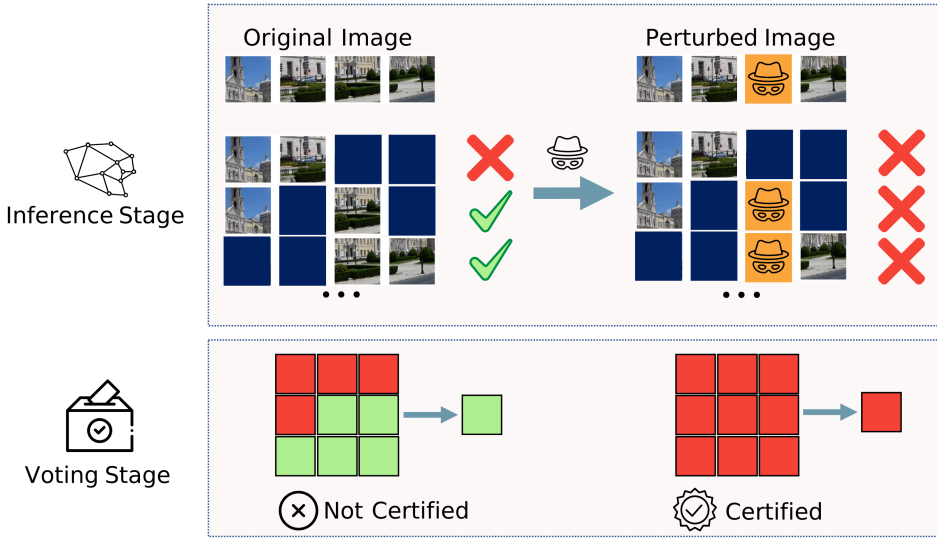


Fig. 14. The illustration that certified detection may fail at test time—the method may return “certified” for attacked images.

neighbourhood is defined as L_p norm in a digital adversarial attack and is defined as a restricted size patch in an adversarial patch attack. To certify such robustness, certified recovery needs to design complicated strategies to ensure the model will yield the same results given only the local features. However, PatchCensor, as well as other certified detection methods, only attempt to detect the presence of abnormal input by making sure that at least one benign voter can control the inference. The philosophy of certified detection can be summarized as follows:

Truth always rests with the minority, and the minority is always stronger than the majority.

However, **What if the minority is wrong?**

This is when the special case comes and why a test-time robustness guarantee can not be given for PatchCensor and all other certified detection work so far. An illustrative example is shown in Figure 14. When the majority voting result is the *correct class* while the other voting support one other class (denoted as class B) unanimously, the prediction is correct. However, the certified detection method will return a “not certified” result as the voting does not reach a consensus. After the attackers perturb the image, the majority voting can all flip to class B. With the original wrong voting, a consensus is reached, and the certified detection method returns a “certified” result by mistake. This subtle difference is ignored in most existing certified detection approaches [30, 54, 89] as the base classifier is assumed to be powerful enough to avoid such cases. For example, in Minority Reports [54], the authors said “...in a benign image, we expect it to be rare for any 3×3 regions in the prediction grid to vote unanimously for an incorrect class.” However, it is indeed possible for the classifier to misclassify an image unanimously. Even powerful vision transformers can make such mistakes on complex datasets like ImageNet, as we find in experiments. An example of such is shown in Figure 15, where the perturbed image is certified as terrapin.

But before concluding that certified detection could be bypassed, we have to answer another question: **when will the minority make mistakes?**

The error made by the benign voters is not controlled by any attacker or natural perturbation, as their influence is completely masked out. It is only because the DL models are imperfect and fail to fit the original distribution well. Moreover, this kind of error can be estimated by measuring




Fig. 15. A real example that can crack the certified adversarial patch detection. The original image has 69 votes for class 35 (mud turtle) and 53 votes for class 36 (terrapin). The perturbed image has 122 votes for class 36. The ground truth of this image is class 35. As a result, the model is correct but not certified on the clean image and wrong but certified on the perturbed image. However, it is hard to distinguish mud turtles from terrapin even for non-expert humans.

Table 13. Adding Additional Information Can Improve the Performance of PatchSmooth

	Single Channel	Three Channel
Clean Accuracy	96.66%	100%
Certified Image Proportion	52.84%	86.03%
$Acc_{in-robust}$	99.9%	100%

when the detection method can return a certified and correct result on the clean dataset. This measurement can provide us with a statistical certificate.

Another important fact is that in certified recovery techniques, “certified robust” also does not mean “certified to be correct,” i.e., the images that pass the certification may also be incorrect due to the imprecision of the models. From this perspective, although the robustness property (i.e., model prediction does not change in the neighbourhood of one input) in common sense can be certified, the guarantee for the model to be correct also comes with uncertainty. Actually, we can observe that the $acc_{in-robust}$ of PatchCensor is always higher than the other two SOTA methods in Figure 9. It is also worth mentioning that in the same experiment, when the ROI is small, there is also a quick drop in accuracy for both certified recovery techniques.

To sum up, although the certification of PatchCensor (and other certified detection techniques) is a statistical guarantee for a data distribution rather than a test-time guarantee for a certain sample, we believe it is more useful in practice due to its superior performance.

6.2 A Possible Way to Improve Certified Recovery

Certified recovery methods aim to use only local features to give reliable predictions. As we have shown in Sections 3 and 5, this can significantly influence their certified accuracy. However, additional features such as texture, the material and the object’s color may help the classifier predict correctly even if only local features are provided. To confirm this, we additionally conduct an experiment on a toy example.

MNIST is itself a single-channel image. We add two channels so that those grayscale images become RGB images. We give different digits with different colors and test the performance of PatchSmooth. The certification is done with band size = 2 and the adversarial patch size = 5. The result is shown in Table 13.




Fig. 16. Our masking strategy works on abstract feature level and can be applied to both images and text.

Although this is a toy example, it still demonstrates one possibility to improve the performance of certified recovery methods. In real-world applications, it is possible to improve the robustness of the camera-based autonomous driving system by diversifying the texture of the traffic signs or adding other additional information. Securely adding more information to local patches could be a potential direction for future certified recovery research.

7 THREATS TO VALIDITY

We summarize the threats to the validity of our study as follows:

Internal validity. The major threat to validity is whether the proposed method can detect abnormal cases and provide a rigorous guarantee. To mitigate this threat, we carefully develop our mutation strategy and prove its correctness in Section 4. Through our evaluation, we also find a possible scenario where the detection could fail and give a detailed discussion of the influence in Section 6. Another threat to validity lies in the training/fine-tuning configuration under different experiment settings. These settings can potentially influence the final performance, and it may be unfair to compare different methods if set improperly. To reduce this threat, we refer to the original paper for each method to find the best hyperparameters for training if they are reported. We retrain/fine-tune the model under different experiment settings when it is necessary.

Construct validity. The construct threats mainly lie in the randomness inherent in our experiments. Specifically, the latency measurement in Section 5 can be unreliable if only a limited number of instances are tested. To reduce such threats, we conduct our experiments by first resizing all the evaluation instances into the same size and recording the latency of each instance. The final result is given by computing the average of the latency.

External validity. The evaluation dataset could be another threat to generalizability. To mitigate this threat, we evaluate our method on the popular CIFAR-10 and the more complex ImageNet dataset. These two classical evaluation datasets can serve as indicators of how powerful our method is. In addition, to validate PatchCensor can achieve moderate performance when ROI is small, we test PatchCensor on PartImageNet and ImageNet object detection, which is closer to real scenarios than the rescaled CIFAR-10.

8 CONCLUSION AND FUTURE WORK

In this article, we propose a simple yet effective certified defense against adversarial patches, which utilizes the novel connection between adversarial patches and input patches in Vision Transformers. We demonstrate the effectiveness and practicality of the defense on ImageNet and CIFAR-10, as well as its flexibility to support different sizes of adversarial patches.

There are two possible directions for future work. First, we can extend PatchCensor for NLP tasks with some adaptation as our masking strategy works on the abstract feature level, and it does not matter whether the input is image or text, as shown in Figure 16. Another direction is

to study the mutation strategy for multiple abnormal patches. The current form of this defense is designed for a single adversarial patch, while other shapes of patches and multiple patches are not handled yet. A problem with multiple patches might be that the number of necessary mutations may be too large, to ensure at least one mutation can exclude the adversarial patches. This problem can be solved by using more fine-grained attention masking strategies and/or probabilistic voting mechanisms, which we leave for future work.

ACKNOWLEDGMENTS

We appreciate Michael McCoyd and David Wagner for the enlightening discussion about the advantages and disadvantages of certified detection defense [54].

REFERENCES

- [1] Ahmed Abusnaina, Yuhang Wu, Sunpreet Arora, Yizhen Wang, Fei Wang, Hao Yang, and David Mohaisen. 2021. Adversarial example detection using latent neighborhood graph. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7687–7696.
- [2] Accompanying Website of this Article. 2022. Retrieved from <https://sites.google.com/view/patchcensor>.
- [3] Andra Acsintoae, Andrei Florescu, Mariana-Iuliana Georgescu, Tudor Mare, Paul Sumedrea, Radu Tudor Ionescu, Fahad Shahbaz Khan, and Mubarak Shah. 2022. UBnormal: New benchmark for supervised open-set video anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20143–20153.
- [4] Yutong Bai, Jieru Mei, Alan L. Yuille, and Cihang Xie. 2021. Are transformers more robust than CNNs? *Adv. Neural Info. Process. Syst.* 34 (2021).
- [5] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. 2021. Understanding robustness of transformers for image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10231–10241.
- [6] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101—mining discriminative components with random forests. In *Proceedings of the European Conference on Computer Vision*. Springer, 446–461.
- [7] Wieland Brendel and Matthias Bethge. 2019. Approximating CNNs with bag-of-local-features models works surprisingly well on ImageNet. In *Proceedings of the 7th International Conference on Learning Representations (ICLR'19)*.
- [8] Tom B. Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. 2017. Adversarial patch. In *Proceedings of the Conference on Neural Information Processing Systems Workshops (NeurIPS Workshops'17)*.
- [9] Nicholas Carlini and David A. Wagner. 2017. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security (AISec@CCS'17)*. 3–14.
- [10] Fabio Carrara, Rudy Becarelli, Roberto Caldelli, Fabrizio Falchi, and Giuseppe Amato. 2018. Adversarial examples detection in features distance spaces. In *Proceedings of the European Conference on Computer Vision (ECCV'18) Workshops*. 0–0.
- [11] Miguel Castro, Barbara Liskov, et al. 1999. Practical byzantine fault tolerance. In *Proceedings of the Operating Systems Design and Implementation Conference (OSDI'99)*. 173–186.
- [12] Zitao Chen, Pritam Dash, and Karthik Pattabiraman. 2021. Turning your strength against you: Detecting and mitigating robust and universal adversarial patch attack. Retrieved from <https://arXiv:2108.05075>.
- [13] Ping-Yeh Chiang, Renkun Ni, Ahmed Abdelkader, Chen Zhu, Christoph Studor, and Tom Goldstein. 2020. Certified defenses for adversarial patches. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*.
- [14] Edward Chou, Florian Tramer, and Giancarlo Pellegrino. 2020. Sentinet: Detecting localized universal attacks against deep learning systems. In *Proceedings of the IEEE Security and Privacy Workshops (SPW'20)*. IEEE, 48–54.
- [15] Dahjung Chung, Khalid Tahboub, and Edward J. Delp. 2017. A two stream siamese convolutional neural network for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision*. 1983–1991.
- [16] Jeremy M. Cohen, Elan Rosenfeld, and J. Zico Kolter. 2019. Certified adversarial robustness via randomized smoothing. In *Proceedings of the 36th International Conference on Machine Learning (ICML'19)*. 1310–1320.
- [17] Patrick Cousot and Radhia Cousot. 1977. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. 238–252.
- [18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly et al. 2021. An image is worth 16×16 words: Transformers for image recognition at scale. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*. Retrieved from <https://openreview.net/forum?id=YicbFdNTTy>.

- [19] Xiaoning Du, Xiaofei Xie, Yi Li, Lei Ma, Yang Liu, and Jianjun Zhao. 2019. Deepstellar: Model-based quantitative analysis of stateful deep learning systems. In *Proceedings of the 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 477–487.
- [20] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. 2018. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'18)*. 1625–1634.
- [21] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. 2018. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1625–1634.
- [22] Yonggan Fu, Shun Yao Zhang, Shang Wu, Cheng Wan, and Yingyan Lin. 2021. Patch-Fool: Are vision transformers always robust against adversarial perturbations? In *Proceedings of the International Conference on Learning Representations*.
- [23] Yonatan Geifman and Ran El-Yaniv. 2019. Selectivenet: A deep neural network with an integrated reject option. In *Proceedings of the International Conference on Machine Learning*. PMLR, 2151–2159.
- [24] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. 2018. ImageNet-trained CNNs are biased toward texture; increasing shape bias improves accuracy and robustness. Retrieved from <https://arXiv:1811.12231>.
- [25] Zhitao Gong, Wenlu Wang, and Wei-Shinn Ku. 2017. Adversarial and clean data are not twins. Retrieved from <https://arXiv:1704.04960>.
- [26] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR'15)*.
- [27] Kathrin Grosse, Praveen Manoharan, Nicolas Papernot, Michael Backes, and Patrick McDaniel. 2017. On the (statistical) detection of adversarial examples. Retrieved from <https://arXiv:1702.06280>.
- [28] Jindong Gu, Volker Tresp, and Yao Qin. 2021. Are vision transformers robust to patch perturbations? Retrieved from <https://arXiv:2111.10659>.
- [29] Jindong Gu, Volker Tresp, and Yao Qin. 2022. Evaluating model robustness to patch perturbations. In *Proceedings of the International Conference on Machine Learning: Shift Happens Workshop (ICML'22)*.
- [30] Husheng Han, Kaidi Xu, Xing Hu, Xiaobing Chen, Ling Liang, Zidong Du, Qi Guo, Yanzhi Wang, and Yunji Chen. 2021. ScaleCert: Scalable certified defense against adversarial patches with sparse superficial layers. *Adv. Neural Info. Process. Syst.* 34 (2021).
- [31] Ju He, Shuo Yang, Shaokang Yang, Adam Kortylewski, Xiaoding Yuan, Jie-Neng Chen, Shuai Liu, Cheng Yang, and Alan Yuille. 2021. PartImageNet: A large, high-quality dataset of parts. Retrieved from <https://arXiv:2112.00933>.
- [32] Dan Hendrycks and Thomas Dietterich. 2018. Benchmarking neural network robustness to common corruptions and perturbations. In *Proceedings of the International Conference on Learning Representations*.
- [33] Wei Huang, Youcheng Sun, Xingyu Zhao, James Sharp, Wenjie Ruan, Jie Meng, and Xiaowei Huang. 2021. Coverage-guided testing for recurrent neural networks. *IEEE Trans. Reliabil.* (2021).
- [34] Nargiz Humbatova, Gunel Jahangirova, and Paolo Tonella. 2021. DeepCrime: Mutation testing of deep learning systems based on real faults. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 67–78.
- [35] Danny Karmon, Daniel Zoran, and Yoav Goldberg. 2018. LaVAN: Localized and visible adversarial noise. In *Proceedings of the 35th International Conference on Machine Learning (ICML'18)*. 2512–2520.
- [36] Kenji Kawaguchi, Leslie Pack Kaelbling, and Yoshua Bengio. 2017. Generalization in deep learning. Retrieved from <https://arXiv:1710.05468>.
- [37] Jinhan Kim, Robert Feldt, and Shin Yoo. 2019. Guiding deep learning system testing using surprise adequacy. In *Proceedings of the IEEE/ACM 41st International Conference on Software Engineering (ICSE'19)*. IEEE, 1039–1049.
- [38] Stepan Komkov and Aleksandr Petiushko. 2021. AdvHat: Real-world adversarial attack on arcfac face id system. In *Proceedings of the 25th International Conference on Pattern Recognition (ICPR'21)*. IEEE, 819–826.
- [39] Zelun Kong, Junfeng Guo, Ang Li, and Cong Liu. 2020. PhysGAN: Generating physical-world-resilient adversarial examples for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14254–14263.
- [40] Alex Krizhevsky. 2009. Learning multiple layers of features from tiny images. Retrieved from <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
- [41] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS'12)*. 1106–1114.
- [42] Mathias Lécuyer, Vaggelis Altidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. 2019. Certified robustness to adversarial examples with differential privacy. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P'19)*. 656–672.

- [43] Alexander Levine and Soheil Feizi. 2020. (De)Randomized smoothing for certifiable defense against patch attacks. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS'20)*.
- [44] Zenan Li, Xiaoxing Ma, Chang Xu, Jingwei Xu, Chun Cao, and Jian Lü. 2020. Operational calibration: Debugging confidence errors for DNNs in the field. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 901–913.
- [45] Ninghao Liu, Hongxia Yang, and Xia Hu. 2018. Adversarial detection with model interpretation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1803–1811.
- [46] Zixi Liu, Yang Feng, and Zhenyu Chen. 2021. DialTest: Automated testing for recurrent-neural-network-driven dialogue systems. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 115–126.
- [47] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. Retrieved from <https://arXiv:2103.14030>.
- [48] Jiajun Lu, Hussein Sibai, Evan Fabry, and David Forsyth. 2017. No need to worry about adversarial examples in object detection in autonomous vehicles. Retrieved from <https://arXiv:1707.03501>.
- [49] Lei Ma, Felix Juefei-Xu, Minhui Xue, Bo Li, Li Li, Yang Liu, and Jianjun Zhao. 2019. Deepct: Tomographic combinatorial testing for deep learning systems. In *Proceedings of the IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER'19)*. IEEE, 614–618.
- [50] Lei Ma, Felix Juefei-Xu, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Chunyang Chen, Ting Su, Li Li, Yang Liu et al. 2018. DeepGauge: Multi-granularity testing criteria for deep learning systems. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. 120–131.
- [51] Lei Ma, Fuyuan Zhang, Jiyuan Sun, Minhui Xue, Bo Li, Felix Juefei-Xu, Chao Xie, Li Li, Yang Liu, Jianjun Zhao et al. 2018. Deepmutation: Mutation testing of deep learning systems. In *Proceedings of the IEEE 29th International Symposium on Software Reliability Engineering (ISSRE'18)*. IEEE, 100–111.
- [52] Wei Ma, Mike Papadakis, Anestis Tsakmalis, Maxime Cordy, and Yves Le Traon. 2021. Test selection for deep learning systems. *ACM Trans. Softw. Eng. Methodol.* 30, 2 (2021), 1–22.
- [53] Kaleel Mahmood, Rigel Mahmood, and Marten Van Dijk. 2021. On the robustness of vision transformers to adversarial examples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7838–7847.
- [54] Michael McCoyd, Won Park, Steven Chen, Neil Shah, Ryan Roggenkemper, Minjune Hwang, Jason Xinyu Liu, and David A. Wagner. 2020. Minority reports defense: Defending against adversarial patches. In *Proceedings of the Applied Cryptography and Network Security Workshops (ACNS Workshops'20)*, Vol. 12418. Springer, 564–582.
- [55] Dongyu Meng and Hao Chen. 2017. MagNet: A two-pronged defense against adversarial examples. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS'17)*. 135–147.
- [56] Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. 2017. On detecting adversarial perturbations. In *Proceedings of the 5th International Conference on Learning Representations (ICLR'17)*.
- [57] Jan Hendrik Metzen and Maksym Yatsura. 2021. Efficient certified defenses against patch attacks on image classifiers. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*. Retrieved from <https://openreview.net/forum?id=hr-3PMvDpil>.
- [58] Mazda Moayeri, Kiarash Banihashem, and Soheil Feizi. 2022. Explicit tradeoffs between adversarial and natural distributional robustness. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS'22)*.
- [59] Muzammal Naseer, Salman Khan, and Fatih Porikli. 2019. Local gradients smoothing: Defense against localized adversarial attacks. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV'19)*. 1300–1307.
- [60] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H. Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. 2021. Intriguing properties of vision transformers. *Adv. Neural Info. Process. Syst.* 34 (2021).
- [61] Nicolas Papernot, Patrick McDaniel, Arunesh Sinha, and Michael P. Wellman. 2018. SoK: Security and privacy in machine learning. In *Proceedings of the IEEE European Symposium on Security and Privacy (EuroS&P'18)*. 399–414.
- [62] Sayak Paul and Pin-Yu Chen. 2021. Vision transformers are robust learners. Retrieved from <https://arXiv:2105.07581>.
- [63] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. 2017. Deepxplore: Automated whitebox testing of deep learning systems. In *Proceedings of the 26th Symposium on Operating Systems Principles*. 1–18.
- [64] Sukrut Rao, David Stutz, and Bernt Schiele. 2020. Adversarial training against location-optimized adversarial patches. In *Proceedings of the European Conference on Computer Vision Workshops (ECCV Workshops'20)*.
- [65] Vincenzo Riccio, Nargiz Humbatova, Gunel Jahangirova, and Paolo Tonella. 2021. DeepMetis: Augmenting a deep learning test set to increase its mutation score. In *Proceedings of the 36th IEEE/ACM International Conference on Automated Software Engineering (ASE'21)*. IEEE, 355–367.
- [66] Giulio Rossolini, Federico Nesti, Gianluca D'Amico, Saasha Nair, Alessandro Biondi, and Giorgio Buttazzo. 2022. On the real-world adversarial robustness of real-time semantic segmentation models for autonomous driving. Retrieved from <https://arXiv:2201.01850>.

- [67] Hadi Salman, Saachi Jain, Eric Wong, and Aleksander Mądry. 2021. Certified patch robustness via smoothed vision transformers. Retrieved from <https://arXiv:2110.07719>.
- [68] Pouya Samangouei, Maya Kabkab, and Rama Chellappa. 2018. Defense-gan: Protecting classifiers against adversarial attacks using generative models. Retrieved from <https://arXiv:1805.06605>.
- [69] Takami Sato, Junjie Shen, Ningfei Wang, Yunhan Jia, Xue Lin, and Qi Alfred Chen. 2021. Dirty road can attack: Security of deep learning based automated lane centering under Physical-World attack. In *Proceedings of the 30th USENIX Security Symposium (USENIX Security'21)*. 3309–3326.
- [70] Dawn Song, Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Florian Tramèr, Atul Prakash, and Tadayoshi Kohno. 2018. Physical adversarial examples for object detectors. In *Proceedings of the 12th USENIX Workshop on Offensive Technologies (WOOT'18)*.
- [71] Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. 2011. The german traffic sign recognition benchmark: A multi-class classification competition. In *Proceedings of the International Joint Conference on Neural Networks*. IEEE, 1453–1460.
- [72] Andrea Stocco, Michael Weiss, Marco Calzana, and Paolo Tonella. 2020. Misbehaviour prediction for autonomous driving systems. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 359–371.
- [73] Hao Su, Jia Deng, and Li Fei-Fei. 2012. Crowdsourcing annotations for visual object detection. In *Proceedings of the Workshops at the 26th AAAI Conference on Artificial Intelligence*.
- [74] Yi Sun, Xiaogang Wang, and Xiaoou Tang. 2016. Sparsifying neural network connections for face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4856–4864.
- [75] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'15)*. 1–9.
- [76] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR'14)*.
- [77] Simen Thys, Wiebe Van Ranst, and Toon Goedemé. 2019. Fooling automated surveillance cameras: Adversarial patches to attack person detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR Workshops'19)*. 49–55.
- [78] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. 2018. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th International Conference on Software Engineering*. 303–314.
- [79] Vladimir Vapnik. 1999. *The Nature of Statistical Learning Theory*. Springer Science & Business Media.
- [80] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 6000–6010.
- [81] Huiyan Wang, Jingwei Xu, Chang Xu, Xiaoxing Ma, and Jian Lu. 2020. Dissector: Input validation for deep learning applications by crossing-layer dissection. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering (ICSE'20)*. IEEE, 727–738.
- [82] Jingyi Wang, Guoliang Dong, Jun Sun, Xinyu Wang, and Peixin Zhang. 2019. Adversarial sample detection for deep neural network through model mutation testing. In *Proceedings of the IEEE/ACM 41st International Conference on Software Engineering (ICSE'19)*. IEEE, 1245–1256.
- [83] Jinwei Wang, Junjie Zhao, Qilin Yin, Xiangyang Luo, Yuhui Zheng, Yun-Qing Shi, and Sunil Kr Jha. 2021. Smsnet: A new deep convolutional neural network model for adversarial example detection. *IEEE Trans. Multimedia* 24 (2021), 230–244.
- [84] Shuai Wang and Zhendong Su. 2020. Metamorphic object insertion for testing object detection systems. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering (ASE'20)*. IEEE, 1053–1065.
- [85] Jian Wei, Jianhua He, Yi Zhou, Kai Chen, Zuoyin Tang, and Zhiliang Xiong. 2019. Enhanced object detection with deep convolutional neural networks for advanced driving assistance. *IEEE Trans. Intell. Transport. Syst.* (2019).
- [86] Ross Wightman. 2019. PyTorch Image Models. Retrieved from <https://github.com/rwightman/pytorch-image-models>. <https://doi.org/10.5281/zenodo.4414861>
- [87] Eric Wong and Zico Kolter. 2018. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *Proceedings of the International Conference on Machine Learning*. PMLR, 5286–5295.
- [88] Chong Xiang, Arjun Nitin Bhagoji, Vikash Sehwal, and Prateek Mittal. 2021. PatchGuard: A provably robust defense against adversarial patches via small receptive fields and masking. In *Proceedings of the 30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association. Retrieved from <https://www.usenix.org/conference/usenixsecurity21/presentation/xiang>.
- [89] Chong Xiang and Prateek Mittal. 2021. PatchGuard++: Efficient provable attack detection against adversarial patches. Retrieved from <https://arXiv:2104.12609>.

- [90] Xiaofei Xie, Lei Ma, Felix Juefei-Xu, Minhui Xue, Hongxu Chen, Yang Liu, Jianjun Zhao, Bo Li, Jianxiong Yin, and Simon See. 2019. Deephunter: A coverage-guided fuzz testing framework for deep neural networks. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 146–157.
- [91] Weilin Xu, David Evans, and Yanjun Qi. 2018. Feature squeezing: Detecting adversarial examples in deep neural networks. In *Proceedings of the 25th Annual Network and Distributed System Security Symposium (NDSS'18)*.
- [92] Puyudi Yang, Jianbo Chen, Cho-Jui Hsieh, Jane-Ling Wang, and Michael Jordan. 2020. MI-loo: Detecting adversarial examples with feature attribution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 6639–6647.
- [93] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M. Ni, and Heung-Yeung Shum. 2022. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. Retrieved from <https://arXiv:2203.03605>.
- [94] Mengshi Zhang, Yuqun Zhang, Lingming Zhang, Cong Liu, and Sarfraz Khurshid. 2018. DeepRoad: GAN-based metamorphic testing and input validation framework for autonomous driving systems. In *Proceedings of the 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE'18)*. IEEE, 132–142.
- [95] Xiyue Zhang. 2020. Uncertainty-guided testing and robustness enhancement for deep learning systems. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering: Companion Proceedings (ICSE-Companion'20)*. IEEE, 101–103.
- [96] Xiyue Zhang, Xiaofei Xie, Lei Ma, Xiaoning Du, Qiang Hu, Yang Liu, Jianjun Zhao, and Meng Sun. 2020. Toward characterizing adversarial defects of deep learning software from the lens of uncertainty. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering (ICSE'20)*. IEEE, 739–751.
- [97] Ziqi Zhang, Yuanchun Li, Yao Guo, Xiangqun Chen, and Yunxin Liu. 2020. Dynamic slicing for deep neural networks. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 838–850.
- [98] Zhishuai Zhang, Cihang Xie, Jianyu Wang, Lingxi Xie, and Alan L. Yuille. 2018. Deepvoting: A robust and explainable deep network for semantic part detection under partial occlusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1372–1380.
- [99] Zhanyuan Zhang, Benson Yuan, Michael McCoyd, and David Wagner. 2020. Clipped BagNet: Defending against sticker attacks with clipped bag-of-features. In *Proceedings of the 3rd Deep Learning and Security Workshop (DLS)*.
- [100] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip H. S. Torr, and Vladlen Koltun. 2021. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 16259–16268.
- [101] Zhe Zhao, Guangke Chen, Jingyi Wang, Yiwei Yang, Fu Song, and Jun Sun. 2021. Attack as defense: Characterizing adversarial examples using robustness. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 42–55.
- [102] Husheng Zhou, Wei Li, Zelun Kong, Junfeng Guo, Yuqun Zhang, Bei Yu, Lingming Zhang, and Cong Liu. 2020. DeepBillboard: Systematic physical-world testing of autonomous driving systems. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering (ICSE'20)*. IEEE, 347–358.

Received 13 September 2022; revised 14 February 2023; accepted 15 February 2023