8-2022

Verifying neural networks against backdoor attacks

Long Hong PHAM
Singapore Management University, hlpham@smu.edu.sg

Jun SUN
Singapore Management University, junsun@smu.edu.sg

Citation

PHAM, Long Hong and SUN, Jun. Verifying neural networks against backdoor attacks. (2022). *Computer Aided Verification: 34th International Conference, CAV 2022, Haifa, Israel, August 7-10: Proceedings*. 13371, 171-192.

Available at: https://ink.library.smu.edu.sg/sis_research/7279

Verifying Neural Networks Against Backdoor Attacks

Long H. Pham(✉) and Jun Sun

Singapore Management University, Singapore, Singapore
{hlpham,junsun}@smu.edu.sg

Abstract. Neural networks have achieved state-of-the-art performance in solving many problems, including many applications in safety/security-critical systems. Researchers also discovered multiple security issues associated with neural networks. One of them is backdoor attacks, i.e., a neural network may be embedded with a backdoor such that a target output is almost always generated in the presence of a trigger. Existing defense approaches mostly focus on detecting whether a neural network is 'backdoored' based on heuristics, e.g., activation patterns. To the best of our knowledge, the only line of work which certifies the absence of backdoor is based on randomized smoothing, which is known to significantly reduce neural network performance. In this work, we propose an approach to verify whether a given neural network is free of backdoor with a certain level of success rate. Our approach integrates statistical sampling as well as abstract interpretation. The experiment results show that our approach effectively verifies the absence of backdoor or generates backdoor triggers.

1 Introduction

Neural networks gradually become an essential component in many real-life systems, e.g., face recognition [25], medical diagnosis [16], as well as auto-driving car [3]. Many of these systems are safety and security-critical. In other words, it is expected that the neural networks used in these systems should not only operate correctly but also satisfy security requirements, i.e., they must sustain attacks from malicious adversaries.

Researchers have identified multiple ways of attacking neural networks, including adversarial attacks [33], backdoor attacks [12], and so on. Adversarial attacks apply a small perturbation (e.g., modifying few pixels in an image input) to a given input (which is often unrecognizable under human inspection) and cause the neural network to generate a wrong output. To mitigate adversarial attacks, many approaches have been proposed, including robust training [7,22], run-time adversarial sample detection [39], and robustness certification [10]. The most relevant to this work is robustness certification, which aims to verify that a neural network satisfies local robustness, i.e., perturbation within a region (e.g., an L_∞ norm) around an input does not change the output. The problem of local robustness certification has been extensively studied in recent years and many methods and tools have been developed [10,14,15,29–32,40,41].

Backdoor attacks work by embedding a 'backdoor' in the neural network so that the neural network works as expected with normal inputs and outputs a specific target

S. Shoham and Y. Vizel (Eds.): CAV 2022, LNCS 13371, pp. 171–192, 2022.
https://doi.org/10.1007/978-3-031-13185-1_9

output in the presence of a backdoor trigger. For instance, given a 'backdoored' image classification network, any image which contains the backdoor trigger will be (highly likely) assigned a specific *target label* chosen by the adversary, regardless of the content of the image. The backdoor trigger can be embedded either through poisoning the training set [12] or modifying a trained neural network directly [19]. It is easy to see that backdoor attacks raise serious security concerns. For instance, the adversaries may use a trigger-containing (a.k.a. 'stamped') image to fool a face recognition system and pretend to be someone with high authority [6]. Similarly, a stamped image may be used to trick an auto-driving system to misidentify street signs and act hazardously [12].

There are multiple active lines of research related to backdoor attacks, e.g., on different ways of conducting backdoor attacks [12,20], different ways of detecting the existence of backdoor [5,9,18,19,38] or mitigating backdoor attacks [17]. Existing approaches are however not capable of certifying the absence of backdoor. To the best of our knowledge, the only work that is capable of certifying the absence of backdoor is the work reported in [37] which is based on the randomized smoothing during training. Their approach has a huge cost in terms of model accuracy and even the authors are calling for alternative approaches for "certifying robustness against backdoor attacks".

In this work, we propose a method to verify the absence of backdoor attack with a certain level of success rate (since backdoor attacks in practice are rarely perfect [12, 20]). Given a neural network and a constraint on the backdoor trigger (e.g., its size), our method is a combination of statistical sampling and deterministic neural network verification techniques (based on abstract interpretation). If we fail to verify the absence of backdoor (due to over-approximation), an optimization-based method is developed to generate concrete backdoor triggers.

We conduct experiments on multiple neural networks trained to classify images in the MNIST dataset. These networks are trained with different types of activation functions, including ReLU, Sigmoid, and Tanh. We verify the absence of backdoor with different settings. The experiment results show that we can verify most of the benign neural networks. Furthermore, we can successfully generate backdoor triggers for neural networks trained with backdoor attack. A slightly surprising result is that we successfully generate backdoor triggers for some of the supposedly benign networks with a reasonably high success rate.

The remaining of the paper is organized as follows. In Sect. 2, we define our problem. In Sect. 3, we present the details of our approach. We show the experiment results in Sect. 4. Section 5 reviews related work and finally, Sect. 6 concludes.

2 Problem Definition

In the following, our discussion focuses on the image domain, in particular, on image classification neural networks. It should be noted that our approach is not limited to the image domain. In general, an image can be represented as a three-dimensional array with shape (c, h, w) where c is the number of channels (i.e., 1 for grayscale images and 3 for color images); h is the height (i.e., the number of rows); and w is the width (i.e., the number of columns) of the image. Each element in the array is a byte value (i.e., from 0 to 255) representing a feature of the image. When an image is used in a classification task with a neural network, its feature values are typically normalized into floating-point

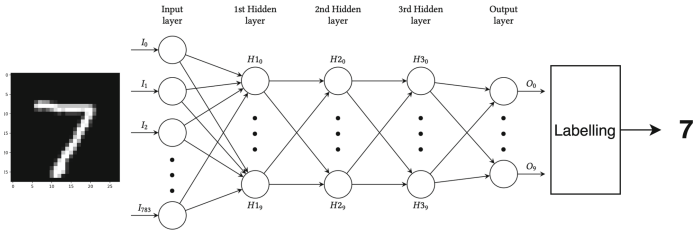

Fig. 1. An example of image classification with neural network

numbers (e.g., dividing the original values by 255 to get normalized values from 0 to 1). Moreover, the image is transformed into a vector with size $m = c \times h \times w$. In this work, we use the three-dimensional form and the vector form of an image interchangeably. The specific form which we use should be clear from the context.

Given a tuple (c_i, h_i, w_i) representing an index in the three-dimensional form, it is easy to compute the according index i in the vector form using the formula: $i = c_i \times h \times w + h_i \times w + w_i$. Similarly, given an index i in the vector form, we compute the tuple (c_i, w_i, h_i) representing the index in the three-dimensional form as follows.

$$
\begin{aligned}
c_i &= i \div (h \times w) \\
h_i &= (i - c_i \times h \times w) \div w \\
w_i &= i - c_i \times h \times w - h_i \times w
\end{aligned}
$$

An image classification task is to label a given image with one of the pre-defined labels automatically. Such tasks are often solved using neural networks. Figure 1 shows the typical workflow of an image classification neural network. The task is to assign a label (i.e., from 0 to 9) to a handwritten digit image. Each input is a grey-scale image with $1 \times 28 \times 28 = 784$ features.

In this work, we focus on fully connected neural networks and convolutional neural networks, which are composed of multiple layers of neurons. The layers include an input layer, a set of hidden layers, and an output layer. The number of neurons in the input layer equals the number of features in the input image. The number of neurons in the output layer equals the number of labels in the classification problem. The number of hidden layers as well as the number of neurons in these layers are flexible. For instance, the network in Fig. 1 has three hidden layers, each of which contains 10 neurons.

The input layer simply applies an identity transformation on the vector of the input image. Each hidden layer transforms its input vector (i.e., the output vector of the previous layer) and produces an output vector for the next layer. Each hidden layer applies two different types of transformations, i.e., the first is an affine transformation and the second is an activation function transformation. Formally, the two transformations of a hidden layer can be defined as: $\vec{y} = \sigma(A * \vec{x} + B)$ where $\vec{x}$ is the input vector, A is the weight matrix, B is the bias vector of the affine transformation, $*$ is the matrix multiplication, σ is the activation function, and $\vec{y}$ is the output vector of the layer. The most popular activation functions include ReLU, Sigmoid, and Tanh. The output layer applies a final affine transformation to its input vector and produces the output vector

(a) Original (b) Stamped

Fig. 2. Some examples of original images and stamped images

of the network. A labelling function $L(\vec{y}) = \arg\max_i \vec{y}$ is then applied on the output vector to return the index of the label with the highest value in $\vec{y}$.

The weights and biases used in the affine transformations are parameters of the neural network. In this work, we focus on pre-trained networks, i.e., the weights and biases of the networks are already fixed. Formally, a neural network is a function $N : R^m \rightarrow R^n = f_k \circ \cdots f_i \cdots \circ f_0$ where m is the number of input features; n is the number of labels; each f_i where $0 < i < k$ is a composition of the affine function and the activation function of the i-th hidden layer; f_0 is the identity transformation of the input layer; and f_k is the last affine transformation of the output layer.

Backdoor Attacks. In [12], Gu *et al.* show that neural networks are subject to backdoor attacks. Intuitively, the idea is that an adversary may introduce a backdoor into the network, for instance, by poisoning the training set. To do that, the adversary starts with choosing a pattern, i.e., a backdoor trigger, and stamps the trigger on a set of samples in the training set (e.g., 20%). Figure 2b shows some stamped images, which are obtained by stamping a trigger to the original images in Fig. 2a. Note that the trigger is a small white square at the top-left corner of the image. A pre-defined target label is the ground truth label for the stamped images. The poisoned training set is then used to train the neural network. The result is a backdoored network that performs normally on clean images (i.e., images without the trigger) but likely assigns the target label to any image which is stamped with the trigger. Besides poisoning the training set, a backdoor can also be introduced by modifying the parameters of a trained neural network directly [19].

Definition 1 (Backdoor trigger). *Given a neural network for classifying images with shape* (c, h, w)*, a backdoor trigger is any image* S *with shape* (c_s, h_s, w_s) *such that* $c_s = c$*,* $h_s \leq h$*, and* $w_s \leq w$*.*

Formally, a backdoor trigger is any stamp that has the same number of channels. Obviously, replacing an input image entirely with a backdoor image with the same size is hardly interesting in practice. Thus, we often limit the size of the trigger. Note that the trigger can be stamped anywhere on the image. In this work, we assume the same trigger is used to attack all images, i.e., the same stamp is stamped at the same position given any input. In other words, we do not consider input-specific triggers, i.e., the triggers that are different for different images. While some forms of input-specific triggers (e.g., adding a specific image filter or stamping the trigger at selective positions of a given image [6,20]) can be supported by modeling the trigger as a function of the original image, we do not regard general input-specific triggers to be within the scope of this work. Given that adversarial attacks can be regarded as a (restricted) form of generating

input-specific triggers, the problem of verifying the absence of input-specific backdoor triggers subsumes the problem of verifying local robustness, and thus the problem is expected to be much more complicated.

Given a trigger with shape (c_s, h_s, w_s), let (h_p, w_p) be the position of the top-left corner of the trigger s.t. $h_p + h_s \leq h$ and $w_p + w_s \leq w$. Given an image I with shape (c, h, w), a backdoor trigger S with shape (c_s, h_s, w_s), and a trigger position (h_p, w_p), a stamped image, denoted as I_s, is defined as follows.

$$I_s[c_i, h_i, w_i] = \begin{cases} S[c_i, h_i - h_p, w_i - w_p] & \text{if } h_p \leq h_i < h_p + h_s \wedge w_p \leq w_i < w_p + w_s \\ I[c_i, w_i, h_i] & \text{otherwise} \end{cases}$$

Intuitively, in the stamped image, the pixels of the stamp replace those corresponding pixels in the original image.

Given a backdoored network, an adversary can perform an attack by feeding an image stamped with the backdoor trigger to the network and expecting the network to classify the stamped image with the target label. Ideally, given any stamped image, an attack on a backdoored network should result in the target label. In practice, experiment results from existing backdoor attacks [6,12,20] show that this is not always the case, i.e., some stamped images may not be classified with the target label. Thus, given a neural network N, a backdoor trigger S, a target label t_s, we say that S has a success rate of θ if and only if there exists a position (h_p, w_p) such that the probability of having $L(N(I_s)) = t_s$ for any I in a chosen test set is θ.

We are now ready to define the problem. *Given a neural network N, a probability of θ and a trigger shape (c_s, h_s, w_s), the problem of verifying the absence of a backdoor attack with a success rate of θ against N is to show that there does not exist a backdoor attack on N which has a success rate of at least θ.*

3 Verifying Backdoor Absence

3.1 Overall Algorithm

The overall approach is shown in Algorithm 1. The inputs include the network N, the required success rate θ, a parameter K representing the sampling size, the trigger shape (c_s, h_s, w_s), the target label t_s, as well as multiple parameters for hypothesis testing (i.e., a type I error α, a type II error β, and a half-width of the indifference region δ). The idea is to apply hypothesis testing, i.e., the SPRT algorithm [1], with the following two mutually exclusive hypotheses.

- H_0: The probability of not having an attack on a set of K randomly selected images is more than $1 - \theta^K$.
- H_1: The probability of not having an attack on a set of K randomly selected images is no more than $1 - \theta^K$.

In the algorithm, variable n and z record the number of times a set of K random images is sampled and is shown to be free of a backdoor with a 100% success rate respectively. Note that function $verifyX$ returns SAFE only if there is no backdoor

Algorithm 1: $verifyPr(N, \theta, K, (c_s, h_s, w_s), t_s, \alpha, \beta, \delta)$

```
1  let n ← 0 be the number of times verifyX is called;
2  let z ← 0 be the number of times verifyX returns SAFE;
3  let p0 ← (1 − θ^K) + δ, p1 ← (1 − θ^k) − δ;
4  while true do
5      n ← n + 1;
6      randomly select a set of images X with size K;
7      if verifyX(N, X, (c_s, h_s, w_s), t_s) returns SAFE then
8          z ← z + 1;
9      else if verifyX(N, X, (c_s, h_s, w_s), t_s) returns UNSAFE then
10         if the generated trigger satisfies the success rate then
11             return UNSAFE;
12     if (p1^z / p0^z) × ((1−p1)^{n−z} / (1−p0)^{n−z}) ≤ β/(1−α) then
13         return SAFE; // Accept H0
14     else if (p1^z / p0^z) × ((1−p1)^{n−z} / (1−p0)^{n−z}) ≥ (1−β)/α then
15         return UNKNOWN; // Accept H1
```

attack on a set of given images X with 100% success rate, i.e., $L(N(I_s)) = t_s$ for all $I \in X$. It may also return a concrete trigger which successfully attacks every image in X. The details of algorithm $verifyX$ is presented in Sect. 3.2.

The loop from lines 4 to 15 in Algorithm 1 keeps randomly selecting and verifying a set of K images using algorithm $verifyX$ until one of the two hypotheses is accepted according to the criteria set by the parameters α and β based on the SPRT algorithm. Furthermore, whenever a trigger is returned by algorithm $verifyX$ at line 9, we check whether the trigger reaches the required success rate on the test set, and return UNSAFE if it does. Note that when H_0 is accepted, we return SAFE, i.e., we successfully verify the absence of a backdoor attack with a success rate of at least θ. When H_1 is accepted, we return UNKNOWN.

Apart from the success rate θ and parameters for hypothesis testing, Algorithm 1 has a particularly interesting parameter K, i.e., the number of images to draw at random each time. On the one hand, if K is set to be small, such as 1, it is very likely algorithm $verifyX$ invoked at line 9 will return UNSAFE since it is often possible to attack a small set of images as demonstrated by many adversarial attack methods [4,11,24], i.e., changing a few pixels of an image changes the output of a neural network. As a result, hypothesis H_1 is accepted and nothing can be concluded. On the other hand, if K is set to be large, such as 10000, due to the complexity of algorithm $verifyX$ (see Sect. 3.2), it is likely that it will timeout and thus return UNKNOWN, which leads to inclusion as well. Furthermore, when K is large, $1 - \theta^K$ will be close to 1 and, as a result, many rounds are needed to accept H_0 even if algorithm $verifyX$ returns SAFE. It is thus important to find an effective K value to balance the two aspects. We identify the value of K empirically in Sect. 4 and aim to study the problem in the future.

Take as an example the network shown in Fig. 1 which is a feed-forward neural network built with the ReLU activation function and three hidden layers. We aim to

verify the absence of a backdoor attack with a success rate of 0.9. We take 10000 images of the MNIST test set to evaluate the success rate of a trigger. We set the parameters in Algorithm 1 as follows: $K = 5$ and $\alpha = \beta = \delta = 0.01$. For the target label 0, after 95 rounds, we have enough evidence to accept the hypothesis H_0, which means we have evidence that there is no backdoor attack on the network with the target label 0 and a success rate of at least 0.9. We have similar results for other target labels, although more rounds of tests are required for labels 2, 3, 5, and 8 (i.e., 98 rounds for label 8, 100 rounds for label 3, 117 rounds for label 5, and 188 rounds for label 2).

3.2 Verifying Backdoor Absence Against a Set of Images

Next, we present the details of algorithm $verifyX$. The inputs include the neural network N, a set of images X with shape (c, h, w), a trigger shape (c_s, h_s, w_s) and a target label t_s. The goal is to check whether exists a trigger which successfully attacks every image in X. Algorithm $verifyX$ may have three outcomes. One is SAFE, i.e., there is no trigger such that backdoor attack succeeds on all the images in X. Another is UNSAFE, i.e., a trigger that can be used to successfully attack all images in X is generated. The last one is UNKNOWN, i.e., we fail to establish either of the above results.

In the following, we describe one concrete realization of the algorithm based on abstract interpretation, as shown in Algorithm 2. At line 1, variable $hasUnknown$ is declared as a flag which is $true$ if and only if we cannot conclude whether there is a successful attack at a certain position. The loop from lines 2 to 15 tries every position for the trigger one by one. Intuitively, variable ϕ is the constraint that must be satisfied by a trigger to successfully attack every image in X. At line 3, we initialize ϕ to be ϕ_{pre}, which is defined as follows: $\phi_{pre} \equiv \bigwedge_{j \in P(h_p, w_p)} lw_j \leq x_j \leq up_j$ where $j \in P(h_p, w_p)$ denotes that j is an index (of an image pixel) in the trigger, x_j is a variable denoting the value of the j-th pixel, lw_j and up_j are the (normalized) minimum (e.g., 0) and maximum (e.g., 1) value of feature j in the image according to the input domain specified by the network N. Intuitively, ϕ_{pre} requires that the pixels in the trigger must be within its domain.

Given a position, the loop from lines 4 to 10 constructs one constraint ϕ_I for each image I, which is the constraint that must be satisfied by the trigger to attack I. In particular, at line 5, function $attackCondition$ is called to construct the constraint. We present the details of this function in Sect. 3.3. If ϕ_I is UNSAT (line 6), attacking image I at position (h_p, w_p) is impossible and we set ϕ to be $false$ and break the loop. Otherwise, we conjunct ϕ with ϕ_I.

After collecting one constraint from each image, we solve ϕ using a constraint solver. If it is not UNSAT (i.e., SAT or UNKNOWN), function $opTrigger$ is called to generate a trigger which is successful on all images in X (if possible). Note that due to over-approximation, the model returned by the solver might be spurious. The details of function $opTrigger$ is presented in Sect. 3.4. If a trigger is successfully generated, we return UNSAFE (at line 13, together with the trigger); otherwise, we set $hasUnknown$ to be $true$ and continue with the next trigger position. Note that we can return UNKNOWN at line 15 without missing any opportunity for verifying the backdoor absence. We instead continue with the next trigger location hoping a trigger may

Algorithm 2: $verifyX(N, X, (c_s, h_s, w_s), t_s)$

```
1  let hasUnknown ← false;
2  foreach trigger position (h_p, w_p) do
3      let φ ← φ_pre;
4      foreach image I ∈ X do
5          let φ_I ← attackCondition(N, I, φ_pre, (c_s, h_s, w_s), (h_p, w_p), t_s);
6          if φ_I is UNSAT then
7              φ ← false;
8              break;
9          else
10             φ ← φ ∧ φ_I;
11     if solving φ results in SAT or UNKNOWN then
12         if opTrigger(N, X, φ, (c_s, h_s, w_s), (h_p, w_p), t_s) returns a trigger then
13             return UNSAFE;
14         else
15             hasUnknown ← true;
16 return hasUnknown ? UNKNOWN : SAFE;
```

be generated successfully. After analyzing all trigger positions (and not finding a successful trigger), if $hasUnknown$ is $true$, we return UNKNOWN or otherwise SAFE.

3.3 Abstract Interpretation

Function $attackCondition$ returns a constraint that must be satisfied such that the trigger with shape (c_s, h_s, w_s) is successful on the image I at position (h_p, w_p). In this work, for efficiency reasons, it is built based on abstract interpretation techniques [32]. Multiple abstract domains have been proposed to analyze neural networks, such as interval [41], Zonotope [30], and DeepPoly [32]. In this work, we adopt the DeepPoly abstract domain [32], which is shown to balance between precision and efficiency.

In the following, we assume each hidden layer in the network is expanded into two separable layers, one for the affine transformation and the other for the activation function transformation. We use l to denote the number of layers in the expanded network, n_i to denote the number of neurons in layer i, and $x^I_{i,j}$ to denote the variable representing the j-th neuron in layer i for the image I. The constraint ϕ_I to be returned by function $attack(N, I, \phi_{pre}, (c_s, h_s, w_s), (h_p, w_p), t_s)$ is a conjunction of three parts.

$$\phi_I \equiv pre_I \wedge \mathcal{A}_{\mathcal{I}} \wedge post_I$$

where pre_I is the constraint on the input features according to the image I, i.e., $pre_I \equiv \phi_{pre} \wedge \left(\bigwedge_{j \in P(h_p,w_p)} x^I_{0,j} = x_j\right) \wedge \left(\bigwedge_{j \notin P(h_p,w_p)} x^I_{0,j} = I[j]\right)$ where $j \notin P(h_p, w_p)$ means that j is not an index (of a pixel) of the trigger; $x^I_{0,j}$ is the variable that represents the input feature j (a.k.a. neuron j at the input layer) of the image I and $I[j]$ is the (normalized) pixel value in the image at index j. Intuitively, the constraint pre_I "erases"

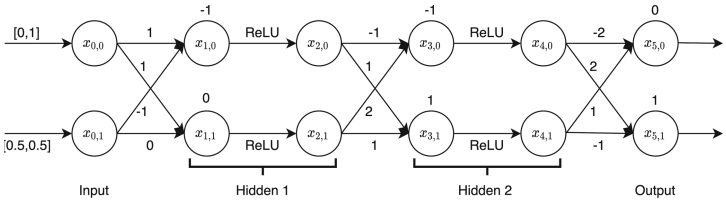

Fig. 3. An example of abstract interpretation

the pixels in the trigger, i.e., they can now take any value with their range, while the remaining pixels must have those value from the image. $post_I$ represents the condition for a successful attack. That is, the value of the target label (i.e., x^I_{l-1,t_s}) must be greater than the values of any other label, i.e., $post_I \equiv \bigwedge_{0 \leq j < n_{l-1} \wedge j \neq t_s} x^I_{l-1,t_s} > x^I_{l-1,j}$.

More interestingly, $\mathcal{A}_{\mathcal{I}}$ is a constraint that over-approximates the behavior of the neural network N according to the DeepPoly abstract domain. That is, given the constraint on the input layer pre_I, a set of abstract transformers are applied to compute a linear over-approximation of every neuron in the next layer, every neuron in the layer after that, and so on until the output layer. The constraint computed on each neuron $x^I_{i,j}$ is of the form $ge^I_{i,j} \leq x^I_{i,j} \leq le^I_{i,j} \wedge lw^I_{i,j} \leq x^I_{i,j} \leq up^I_{i,j}$ where $ge^I_{i,j}$ and $le^I_{i,j}$ are two linear expressions constituted by variables representing neurons from the previous layer (i.e., layer $i-1$); and $lw^I_{i,j}$ and $up^I_{i,j}$ are the concrete lower bound and upper bound of the neuron. Note that the abstract transformers are different for the activation function layer and affine layer. As the DeepPoly abstract transformers are not our contribution, we skip the details and refer the reader to [32] for details on the abstract transformers, including their soundness (i.e., they always over-approximate).

Example 1. Since it is too complicated to show the details of applying abstract interpretation to the neural network shown in Fig. 1, we instead construct a simple example as shown in Fig. 3 to illustrate how it works. There are two features in this artificial image I, i.e., $x^I_{0,1}$ has a constant value of 0.5 and $x^I_{0,0}$ is the trigger whose value ranges from 0 to 1. That is, $pre_I \equiv 0 \leq x^I_{0,0} \leq 1 \wedge x^I_{0,1} = 0.5$. After expanding the hidden layers, the network has 6 layers, each of which has 2 neurons. Applying the DeepPoly abstract transformers from the input layer all the way to the output layer, we obtain the abstract states for the last layer. Further, assume that the target label is 0. The constraint $post_I$ is thus as follows: $post_I \equiv x^I_{5,0} > x^I_{5,1}$. Solving the constraints returns SAT with $x^I_{0,0} = 0$. Indeed, with the stamped image $I_s = [0, 0.5]$, the output vector is $[1, 0]$. We thus identified a successful attack on the target label 0.

Optimization. Note that at line 6 of Algorithm 2, for each constraint ϕ_I, we perform a quick check to see if the constraint is satisfiable or not. If ϕ_I is UNSAT, we can ignore the remaining images and analyze the next trigger position, which allows us to speed up the process. One naive approach is to call a solver on ϕ_I, which would incur significant overhead since it could happen many times. To reduce the overhead, we propose a simple procedure to quickly check whether ϕ_I is UNSAT based solely on its abstract states at the output layer. That is, we check the satisfiability of the following constraint

instead: $\bigwedge_{0 \leq j < n_{l-1} \wedge j \neq t_s} up^I_{l-1,t_s} > lw^I_{l-1,j}$. Recall that up^I_{l-1,t_s} is the concrete upper bound of the neuron t_s and $lw^I_{l-1,j}$ is the concrete lower bound of the neuron j at the output layer. Thus, intuitively, we check whether the concrete upper bound of the target label t_s is larger than the concrete lower bound of every other label. If it is UNSAT, it is impossible to have the target label as the result and thus the attack would fail on the image I. We then only call the solver on ϕ_I if the above procedure does not return UNSAT. Furthermore, the loop in Algorithm 2 can be parallelized straightforwardly, i.e., by using a separate process to verify against a different trigger position. Whenever a trigger is found by any of the processes, the whole algorithm is then interrupted.

3.4 Generating Backdoor Triggers

In the following, we present the details of function $opTrigger$, which intuitively aims to generate a trigger S with shape (c_s, h_s, w_s) at position (h_p, w_p) for attacking every image I in X successfully. If the solver applied to solve ϕ at line 11 of Algorithm 2 returns a model that satisfies ϕ, we first check whether the model is indeed a trigger that successfully attacks every image in X. Due to over-approximation of abstract interpretation, the model might be a spurious trigger. If it is a real trigger, we return the model. Otherwise, we employ an optimization-based approach to generate a trigger.

Given a network N, one image I, a target label t_s, and a position (h_p, w_p), let I_s is the stamped image generated from I by stamping I with the trigger at the position (h_p, w_p). We generate a backdoor trigger S by minimizing the following loss function.

$$loss(N, I, S, (h_p, w_p), t_s) = \begin{cases} 0 & \text{if } n_s > n_o \\ (n_o - n_s + \epsilon) & \text{otherwise} \end{cases}$$

where $n_s = N(I_s)[t_s]$ is the output value of the target label; $n_o = \max_{j \neq t_s} N(I_s)[j]$ is the maximum value of any label other than the target label; and ϵ is a small constant (e.g., 10^{-9}). Note that the trigger S is the only variable in the loss function. Intuitively, the loss function returns 0 if the attack on I by the trigger is successful. Otherwise, it returns a quantitative measure on how far the attack is from being successful on attacking I. Given a set of images X, the loss function is defined as the sum of the loss for each image I in X: $loss(N, X, S, (h_p, w_p), t_s) = \sum_{I \in X} loss(N, I, S, (h_p, w_p), t_s)$. The following optimization problem is then solved to find an attack which successfully attacks all images in X: $\arg\min_S loss(N, X, S, (h_p, w_p), t_s)$.

3.5 Correctness and Complexity

Lemma 1. *Given a neural network N, a set of images X, a trigger shape (c_s, h_s, w_s), and a target label t_s, Algorithm 2 (1) returns SAFE only if there is no backdoor attack which is successful on all images in X with the provided trigger shape and target label; and (2) returns UNSAFE only if there exists a backdoor attack which is successful on all images in X with the provided trigger shape and target label.*

Proof. By [32], function $attackCondition$ always returns a constraint which is an over-approximation of the constraint that must be satisfied such that the trigger is successful on image I. Furthermore, Algorithm 2 returns SAFE only at line 16, i.e., only

if constraints that must be satisfied to attack all images in X at each certain position are UNSAT. Thus, (1) is established. (2) is trivially established since we only return UNSAFE when a trigger that is successful on every provided image is generated. □

The following establishes the soundness of our approach.

Theorem 1. *Given a neural network N, a success rate θ, a target label t_s, a trigger shape (c_s, h_s, w_s), a type I error α, a type II error β, and a half-width of the indifference region δ, Algorithm 1 returns SAFE only if there is sufficient evidence (subject to type I error α and type II error β) that there is no backdoor attack with a success rate at least θ with the provided trigger shape and target label at the specified significance level.*

Proof. If there is a backdoor attack with a success rate no less than θ, given a set of randomly K selected images, the probability of having an attack is no less than θ^K (since there is at least one backdoor attack with a success rate no less than θ and maybe more). Thus, the probability of not having an attack is no more than $1 - \theta^K$. By the correctness of the SPRT algorithm, Algorithm 1 returns SAFE only if there is sufficient evidence that H_0 is true, i.e., the probability of not having an attack on a set of K randomly selected images is more than $1 - \theta^K$, implying it is sufficient evidence that there is no backdoor attack with success rate no less than θ. The theorem holds. □

Furthermore, it is trivial to show that Algorithm 1 returns UNSAFE only if there exists a backdoor attack which has a success rate at least θ with the provided trigger shape and target label.

In the following, we briefly discuss the complexity of our approach. It is straightforward to see that Algorithm 2 always terminates if a timeout is imposed on solving the constraints and the optimization problems. Since we can always set a tight time limit on solving the constraints and the optimization problems, the complexity of the algorithm is determined mainly by the complexity of function $attackCondition$, which in turn is determined by the complexity of abstract interpretation. The complexity of applying abstract interpretation with the DeepPoly abstract domain is $\mathcal{O}(l^2 \times n_{max}^3)$ where l is the number of layers, and n_{max} is the maximum number of neurons in any of the layers. Let K be the number of images in X. Note that the number of trigger positions is $\mathcal{O}(h \times w)$, i.e., the size of an image. The best case complexity of Algorithm 2 is $\mathcal{O}(l^2 \times n_{max}^3 \times h \times w)$ and the worst case complexity is $\mathcal{O}(l^2 \times n_{max}^3 \times K \times h \times w)$. We remark that in practice, l typically ranges from 1 to 20; n_{max} is often advised to be no more than the input size (e.g., from dozens to thousands usually); K ranges from a few to hundreds; and $h \times w$ depends on the image resolution (e.g., from hundreds to millions). Thus, in general, Algorithm 2 could be time-consuming in practice and we anticipate further optimization in future work.

The complexity of Algorithm 1 is the complexity of Algorithm 2 times the complexity of the SPRT algorithm. The complexity of the SPRT algorithm is in general hard to quantify and we refer the readers to [1] for a detailed discussion.

3.6 Discussion

Our approaches are designed to verify the absence of input-agnostic (i.e., not input-specific) backdoor attacks as presented in Sect. 2. In the following, we briefly review other backdoor attacks and discuss how to extend our approach to support them.

In [12], Gu *et al.* described a backdoor attack which, instead of forcing the network to classify any stamped image with the target label, only alters the label if the original image has a specific ground truth label t_i (e.g., Bob with the trigger will activate the backdoor and be classified as Alice the manager). Our verification approach can be easily adapted to verify the absence of this attack by focusing on images with label t_i in Algorithm 1 and Algorithm 2.

Another attack proposed in [12] works by reducing the performance (e.g., accuracy) of the neural network on the images with a specific ground truth label t_i, i.e., given an image with ground truth label t_i, the network will classify the stamped image with some label $t_s \neq t_i$. The attack can be similarly handled by focusing on images with ground truth label t_i, although due to the disjunction introduced by $t_s \neq t_i$, the constraints are likely to be harder to solve. That is, we can focus on images with ground truth label t_i in Algorithm 2, and define an attack to be successful if $L(N(I_s)) \neq t_i$ is satisfied.

In [19], Liu *et al.* proposed to use backdoor triggers with different shapes (i.e., not just in the form of a square or a rectangle). If the user is aware of the shape of the backdoor trigger, a different trigger can be used as input for Algorithm 1 and Algorithm 2 and the algorithms would work to verify the absence of such backdoor. Alternatively, the users can choose a square-shaped backdoor trigger that is larger enough to cover the actual backdoor trigger, in which case our algorithms would remain to be sound, although it might be inconclusive if the trigger is too big.

Multiple groups [2,20,28,35] proposed the idea of poisoning only those samples in the training data which have the same ground truth label as the target label to improve the stealthiness of the backdoor attack. This type of attack is designed to trick the human inspection on the training data, and so does not affect our verification algorithms.

In this work, we consider a specific type of stamping, i.e., the backdoor trigger replaces the part of the original clean image. Multiple groups [6,19] proposed the use of the blending operation as a way of 'stamping', i.e., the features of the backdoor trigger are blended with the features of the original images with some coefficients α. This is a form of input-specific backdoor, the trigger is different for different images. To handle such kind of backdoor attacks, one way is to modify the constraint pre_I according to the blending operation (assuming that α is known). Since the blending operation proposed in [6,19] is linear, we expect this would not introduce additional complexity to our algorithms.

Input-specific triggers, in general, may pose a threat to our approach. First, some input-specific triggers [19,20] cover the whole image, which is likely to make our approach inclusive due to false alarms resulted from over-approximation. Second, it may not be easy to model some of the input-specific triggers in our framework. For instance, Liu *et al.* [20] recently proposed to use reflection to create stamped images that look natural. Modeling the 'stamping' operation for this kind of attack would require us to know where the reflection is in the image, which is highly non-trivial. However, it should also be noted that input-specific triggers are often not as effective as input-agnostic triggers, e.g., the reflection-based attack reported in [20] are hard to reproduce. Furthermore, as discussed in Sect. 2, backdoor attack with input-specific triggers is an attacking method that is more powerful than adversarial attacks, and the problem of verifying the absence of backdoor attack with input-specific triggers is not yet clearly defined.

4 Implementation and Evaluation

We have implemented our approach as a self-contained analysis engine in the Socrates framework [26]. We use Gurobi [13] to solve the constraints and use scipy [36] to solve the optimization problems.

We collect a set of 51 neural networks. 45 of them are fully connected ones and are trained on the MNIST training set (i.e., a standard dataset which contains black and white images of digits). These networks have the number of hidden layers ranging from 3 to 5. For each network, the number of neurons in each of its hidden layers ranges from 10 to 50, i.e., 10, 20, 30, 40, or 50. To evaluate our approach on neural networks built with different activation functions, each activation function (i.e., ReLU, Sigmoid, and Tanh) is used in 15 of the neural networks. Among the remaining six networks, three of them are bigger fully connected networks adopted from the benchmarks reported in [32]. They are all built with the ReLU activation function. For convenience, we name the networks in the form of *f_k_n* where *f* is the name of the activation function, *k* is the number of hidden layers, and *n* is the number of neurons in each hidden layer. The remaining three networks are convolutional networks (which are often used in face recognition systems) adopted from [32]. Although they have the same structure, i.e., each of them has two convolutional hidden layers and one fully connected hidden layer, they are trained differently. One is trained in the normal way; one is trained using DiffAI [22], and the last one is trained using projected gradient descent [7]. These training methods are developed to improve the robustness of neural networks against adversarial attacks. Our aim is thus to evaluate whether they help to prevent backdoor attacks as well. We name these networks *conv*, *conv_diffai*, and *conv_pgd*.

We verify the networks against the backdoor trigger with shape $(1, 3, 3)$. All the networks are trained using clean data since we focus on verifying the absence of backdoor attacks. They all have precision of at least 90%, except *Sigmoid_4_10* and *Sigmoid_5_10*, which have precision of 81% and 89% respectively. In the following, we answer multiple research questions. All the experiments are conducted using a machine with 3.1Ghz 16-core CPU and 64GB RAM. All models and experiment details are at [27].

RQ1: Is our realization of $verifyX$ *effective?* This question is meaningful as our approach relies on Algorithm $verifyX$. To answer this question, for each network, we select the first 100 images in the test set (i.e., a K of 100 for Algorithm 1, which is more than sufficient) and then apply Algorithm $verifyX$ with these images and each of the labels, i.e., 0 to 9. In total, we have 510 verification tasks. For each network, we run 10 processes in parallel, each of which verifies a separate target. The only exception is the network *ReLU_3_1024*, due to its complexity, we only run five parallel processes (since each process consumes a lot of resources). In each verification process, we filter out those images which are classified wrongly by the network as well as the images which are already classified as the target label.

Figure 4 shows the results. The x-axis show the groups of the networks, e.g., *ReLU_3* means five fully connected networks using the ReLU activation function with three hidden layers; *3 Full* and *3 Conv* mean the three fully connected and the three convolutional networks adapted from [32] respectively. The y-axis shows the number of (network, target) pairs. Note that each group may contain a different number of pairs, i.e., the

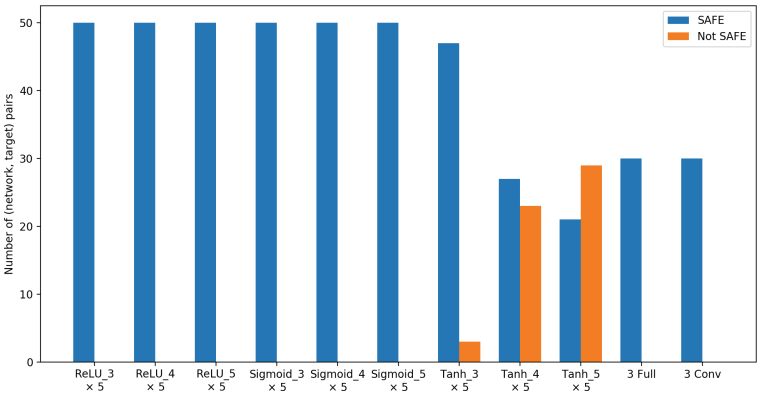

Fig. 4. The results of $verifyX$

maximum values for the small network groups are 50, and the maximum values for the last two groups are 30. First, we successfully verify 455 out of 510 verification tasks (i.e., 89%) of them, i.e., the neural network is safe with respect to the selected images. It is encouraging to notice that the verified tasks include all models adopted from [32], which are considerably larger (e.g., with 1024 neurons at each layer) and more complex (i.e., convolutional networks). Second, some networks are not proved to be safe with some target labels. It could be either there is indeed a backdoor trigger that we fail to identify (through optimization), or we fail to verify due to the over-approximation introduced by abstract interpretation. Lastly, with the same structure (i.e., the same number of hidden layers and the same number of neurons in each hidden layer), the networks using the ReLU and Sigmoid activation functions are more often verified to be safe than those using the Tanh activation function. This is most likely due to the difference in the precision of the abstract transformers for these functions.

RQ2: can we verify the absence of backdoor attacks with a certain level of success rate? To answer this question, we evaluate our approach on six networks used in RQ1, i.e., *ReLU_3_10*, *ReLU_5_50*, *Sigmoid_3_10*, *Sigmoid_5_50*, *Tanh_3_10*, and *Tanh_5_50*. These networks are chosen to cover a wide range of the number of hidden layers and the number of neurons in each layer, as well as different activation functions. Note that due to the high complexity of Algorithm 1 (which potentially applies Algorithm 2 hundreds of times), running Algorithm 1 on all the networks evaluated in RQ1 requires an overwhelming amount of resources. *Furthermore, since there is no existing work on backdoor verification, we do not have any baseline to compare with.*

Recall that Algorithm 1 has two important parameters K and θ, both of which potentially have a significant impact on the verification result. We thus run each network with four different settings, in which the number of images K is set to be either 5 or 10, and the success rate θ is either 0.8 or 0.9. In total, with 10 target labels, we have a total of 240 verification tasks for this experiment. Note that some preliminary experiments are conducted before we select these two K values.

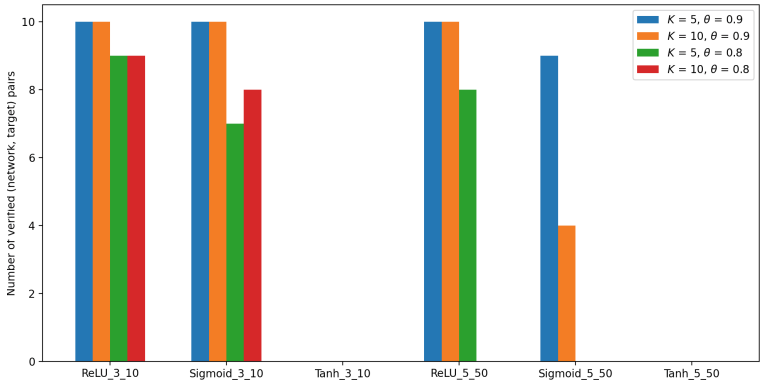

Fig. 5. Verification results

We use all the 10000 images in the test set as the image population and randomly choose K images in each round of test. When a trigger is generated, the success rate of the trigger is validated on the images in the test set (after the above-mentioned filtering). Like in RQ1, we run each network with 10 parallel processes, each of which verifies a separate target. As the SPRT algorithm may take a very long time to terminate, we set a timeout for each verification task, i.e., 2 h for those networks with three hidden layers, and 10 h for those networks with five hidden layers.

The results are shown in Fig. 5. The x-axis shows the networks, the y-axis shows the number of verified pairs of network and target label. We have multiple observations based on the experiment results. First, a quick glance shows that with the same structure and hypothesis testing parameters, more networks built with the ReLU activation function are verified than those built with the Sigmoid and Tanh functions. Second, we notice that the best result is achieved with $K = 5$ and $\theta = 0.9$. With these parameter values, we can verify that three networks *ReLU_3_10*, *ReLU_5_50*, and *Sigmoid_3_10* are safe with respect to all the target labels and the network *Sigmoid_5_50* is safe with respect to nine over 10 target labels. If we keep the same success rate as 0.9 and increase the number of images K from 5 to 10, we can see that the number of verified cases in the network *Sigmoid_5_50* decreases. This is because when we increase the number of images that must be attacked successfully together, the probability that we do not have the attack increases, which means we need more rounds of test to confirm the hypothesis H_0 and so the verification process for the network *Sigmoid_5_50* times out before reaching the conclusion. We have a similar observation when we keep the number of images K at 5 but decrease the success rate from 0.9 to 0.8. When the success rate decreases, the probability of not having the attack increases, which requires more tests to confirm the hypothesis H_0. As a result, for all these four networks, there are multiple verification tasks that time out before reaching the conclusion. However, we notice that there is an exception when we keep the success rate as 0.8 and increase the number of images from 5 to 10. While the number of verified cases for the network *ReLU_5_50* decreases (which can be explained in the same way as above), the number of verified cases for the network *Sigmoid_3_10* increases (and the results for the other two

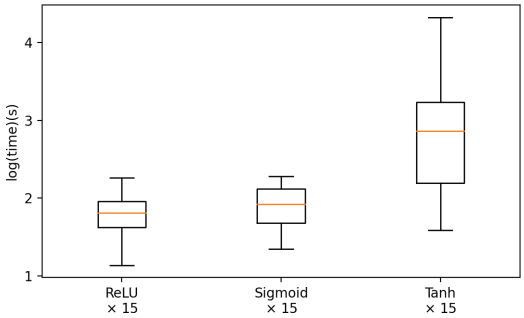

Fig. 6. The running time of the experiments in RQ1 with benchmark networks

networks do not change). Our explanation is that when we increase the number of images K to 10, it is easier for the Algorithm 2 to conclude that there is no attack, and so the Algorithm 1 still collects enough evidence to conclude H_0. On the other hand, with the number of images is 5, Algorithm 2 may return a lot of UNKNOWN (due to spurious triggers), and so the hypothesis testing in the Algorithm 1 goes back and forth between the two hypotheses H_0 and H_1 and eventually times out.

A slightly surprising result is obtained for the network *Tanh_3_10*, i.e., our trigger generation process generates two triggers for the target labels 2 and 5 when the success rate is set to be 0.8. This is surprising as these networks are not generated with backdoor attack. This result can be potentially explained by the combination of the relatively low success rate (i.e., 0.8) and the phenomenon known as universal adversarial perturbations [23]. With the returned triggers, the users may want to investigate the network further and potentially improve it with techniques such as robust training [7,22].

RQ3: Is our approach efficient time-wise? To answer this question, we collect the wall-clock time to run the experiments in RQ1 and RQ2. For each network, we record the average running time for 10 different target labels. The results for 45 small networks are shown in Fig. 6. The x-axis shows the groups of 15 networks categorized based on their activation functions and the y-axis shows the logarithmic scale of the running time in the form of boxplots (where the box shows the result of 25 percentile to 75 percentile, the bottom and top lines are the minimum and maximum, and the orange line is median). The execution time ranges from 14 s to less than 6 h for these networks. Furthermore, we can see that there is not much difference between the running time of the networks using the ReLU and Sigmoid activation functions. However, the running time of the networks using the Tanh function is one order of magnitude larger than those of the ReLU and Sigmoid networks. The reason is that the Tanh networks have many non-safe cases (as shown in Fig. 4) and, as a result, the verification process needs to check more images at more trigger positions. The running time of those networks adopted from [32] ranges from more than 5 min to less than 4 h, as shown in Table 1. Finally, the running time for each network in RQ2 (i.e., the time required to verify the networks against backdoor attacks) according to different settings is shown in Table 2.

Table 1. The running time of the experiments in RQ1 with networks adapted from [32]

Network	Time	Network	Time
ReLU_3_1024	237 m 24s	conv	194 m 30 s
ReLU_5_100	5 m 38 s	conv_diffai	111 m 12 s
ReLU_8_200	48 m 34s	conv_pgd	190 m 19 s

Table 2. The running time of the experiments in RQ2

Network	$K = 5$ $\theta = 0.9$	$K = 10$ $\theta = 0.9$	$K = 5$ $\theta = 0.8$	$K = 10$ $\theta = 0.8$
ReLU_3_10	31 m 31 s	46 m 39 s	55 m 44 s	68 m 54 s
ReLU_5_50	341 m 36 s	493 m 30 s	551 m 40 s	600 m 0 s
Sigmoid_3_10	46 m 43 s	59 m 28 s	92 m 34s	93 m 21 s
Sigmoid_5_50	476 m 38 s	588 m 25 s	600 m 0s	600 m 0 s
Tanh_3_10	114 m 2 s	105 m 18 s	50 m 58 s	26 m 4 s
Tanh_5_50	600 m 0s	600 m 0 s	600 m 0 s	600 m 0 s

RQ4: can our approach generate backdoor triggers? Being able to generate counterexamples is a part of a useful verification method. We conduct another experiment to evaluate the effectiveness of our backdoor trigger generation approach. We train a new set of 45 networks that have the same structure as those used for answering RQ1. The difference is that this time each network is trained to contain backdoor through data poisoning. In particular, for each network, we randomly extract 20% of the training data, stamp a white square with shape $(1, 3, 3)$ in one corner of the images, assign a random target label, and then train the neural network from scratch with the poisoned training data. While such an attack is shown to be effective [12], it is not guaranteed to be always successful on a randomly selected set of images. Thus, we do the following to make sure that there exists a trigger for a set of selected images. From 10000 images in the test set, we first filter out those images which are classified wrongly or already classified with the target label. The remaining images are collected into a set X_0. Next, to make sure that the selected images have a high chance of being attacked successfully, we apply another filter on X_0. This time, we stamp each image in X_0 with a white square at the same trigger position as we poison the training data. We then keep the image if its stamped version is classified by the network with the target label. The remaining images after the second filter are collected into another set X. We apply our approach, in particular, the backdoor trigger generation on X, if $|X| \div |X_0| \geq 0.8$, i.e., the backdoor attack has a success rate of 80%.

The results are shown in Fig. 7 in which the y-axis shows the number of networks. The timeout is set to be 120 s. Among the 45 networks, we can see that a trigger is successfully generated for 33 (i.e., 73%) of the networks. A close investigation of these networks shows that the generated trigger is the exact white square that is used to stamp the training data. There are 12 networks for which the trigger is not generated. We

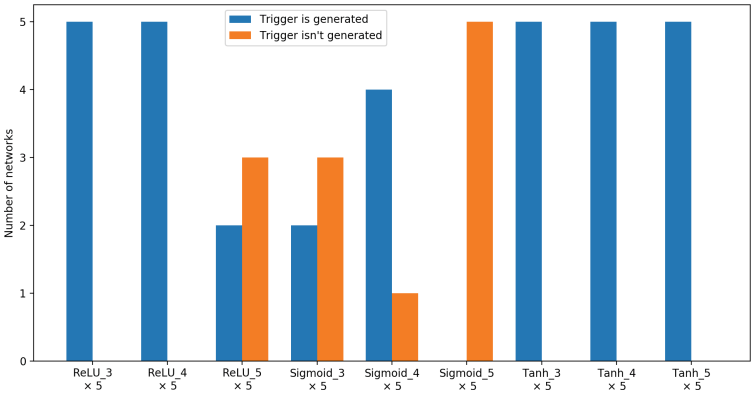

Fig. 7. The results of backdoor trigger generation

investigate these networks and see that they are either too biased (i.e., classifying every image with the target label and thus $|X_0| = 0$) or the attack on these networks does not perform well (i.e., $|X| \div |X_0| < 0.8$). In other words, the backdoor attack on these networks failed and, as a result, the generation process does not even begin with these networks. In a nutshell, we successfully generate the trigger for every successful backdoor attack. Finally, note that the running time of the backdoor generation process is reasonable (i.e., on average, 50 s to generate a backdoor trigger for one network) and thus it does not affect the overall performance of our verification algorithm.

5 Related Work

The work which is closest to ours is [37] in which Wang *et al.* aim to certify neural networks' robustness against backdoor attack using randomized smoothing. However, there are many noticeable differences between their approach and ours. First, while our work focuses on verifying the absence of backdoor, their work aims to certify the robustness of individual images based on the provided training data and learning algorithm (which can be used to implicitly derive the network). Second, by using random noises to estimate the networks' behaviors, their approach can only obtain very loose results. As shown in their experiments, they can only certify the robustness against backdoor attack with triggers contains two pixels and on a "toy" network with only two layers and two labels, after simplifying the input features by rounding them into 0 or 1. Compare to their approach, our approach can apply to networks used to solve real image classification problems as shown in our experiments.

Our work is closely related to a line of work on verifying neural networks. Existing approaches mostly focus on local robustness property and can be roughly classified into two categories: exact methods and approximation methods. The exact methods aim to model the networks precisely and solve the verification problem using techniques such as mixed-integer linear programming [34] or SMT solving [8,15]. On the one hand, these approaches can guarantee sound and complete results in verifying neural networks. On the other hand, they often have limited scalability and thus are limited to

small neural networks. Moreover, these approaches have difficulty in handling activation functions except the ReLU function.

In comparison, the approximation approaches over-approximate neural network behavior to gain better scalability. AI^2 [10] is the first work pursuing this direction using the classic abstract interpretation technique. After that, more researchers try to explore different abstract domains for better precision without sacrificing too much scalability [29,30,32]. In general, the approximation approaches are more scalable than the exact methods, and they are capable of handling activation functions such as Sigmoid and Tanh. However, due to the over-approximation, these methods may fail to verify a valid property.

We also notice that it is possible to incorporate abstraction refinement to the approximation methods and gain better precision, for instance, by splitting an abstraction into multiple parts to reduce the imprecision due to over-approximation. There are many works [21,40,41] which fall into this category. We remark that our approach is orthogonal to the development of sophisticated verification techniques for neural networks.

Finally, our approach, especially the part on backdoor trigger generation, is related to many approaches on generating adversarial samples for neural networks. Some representative approaches in this category are FGSM [11], JSMA [24], and C&W [4] which aim to generate adversarial samples to violate the local robustness property, and [42] which aims to violate fairness property.

6 Conclusion

In this work, we propose the first approach to formally verify that a neural network is safe from backdoor attacks. We address the problem on how to verify the absence of a backdoor that reaches a certain level of success rate. Our approach is based on abstract interpretation and we provide an implementation based on DeepPoly abstract domain. The experiment results show the potential of our approach. In the future, we intend to extend our approach with more abstract domains as well as improve the performance to verify more real-life networks. Besides that, we also intend to apply our approach to verify the networks designed for other tasks, such as sound or text classification.

Acknowledgements. This research is supported by the Ministry of Education, Singapore under its Academic Research Fund Tier 3 (Award ID: MOET32020-0004). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore. This research is also partly supported by the Starry Night Science Fund of Zhejiang University Shanghai Institute for Advanced Study, Grant No. SN-ZJU-SIAS-001.

References

1. Agha, G., Palmskog, K.: A survey of statistical model checking. ACM Trans. Model. Comput. Simul. (TOMACS) **28**(1), 1–39 (2018)
2. Barni, M., Kallas, K., Tondi, B.: A new backdoor attack in CNNs by training set corruption without label poisoning. In: ICIP 2019, pp. 101–105. IEEE (2019)

3. Bojarski, M., et al.: End to end learning for self-driving cars. arXiv preprint arXiv:1604.07316 (2016)
4. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: S&P 2017, pp. 39–57. IEEE (2017)
5. Chen, B., et al.: Detecting backdoor attacks on deep neural networks by activation clustering. arXiv preprint arXiv:1811.03728 (2018)
6. Chen, X., Liu, C., Li, B., Lu, K., Song, D.: Targeted backdoor attacks on deep learning systems using data poisoning. arXiv preprint arXiv:1712.05526 (2017)
7. Dong, Y., et al.: Boosting adversarial attacks with momentum. In: CVPR 2018, pp. 9185–9193. IEEE (2018)
8. Ehlers, R.: Formal verification of piece-wise linear feed-forward neural networks. In: D'Souza, D., Narayan Kumar, K. (eds.) ATVA 2017. LNCS, vol. 10482, pp. 269–286. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-68167-2_19
9. Gao, Y., Xu, C., Wang, D., Chen, S., Ranasinghe, D.C., Nepal, S.: Strip: a defence against trojan attacks on deep neural networks. In: ACSAC 2019, pp. 113–125. ACM (2019)
10. Gehr, T., Mirman, M., Drachsler-Cohen, D., Tsankov, P., Chaudhuri, S., Vechev, M.: Ai2: Safety and robustness certification of neural networks with abstract interpretation. In: S&P 2018, pp. 3–18. IEEE (2018)
11. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014)
12. Gu, T., Dolan-Gavitt, B., Garg, S.: Badnets: Identifying vulnerabilities in the machine learning model supply chain. arXiv preprint arXiv:1708.06733 (2017)
13. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2021). https://www.gurobi.com
14. Huang, X., Kwiatkowska, M., Wang, S., Wu, M.: Safety Verification of Deep Neural Networks. In: Majumdar, R., Kunčak, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 3–29. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63387-9_1
15. Katz, G., Barrett, C., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: An efficient smt solver for verifying deep neural networks. In: CAV 2017, pp. 97–117. Springer (2017)
16. Li, Q., Cai, W., Wang, X., Zhou, Y., Feng, D.D., Chen, M.: Medical image classification with convolutional neural network. In: ICARCV 2014, pp. 844–848. IEEE (2014)
17. Liu, K., Dolan-Gavitt, B., Garg, S.: Fine-pruning: defending against backdooring attacks on deep neural networks. In: Bailey, M., Holz, T., Stamatogiannakis, M., Ioannidis, S. (eds.) RAID 2018. LNCS, vol. 11050, pp. 273–294. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00470-5_13
18. Liu, Y., Lee, W.C., Tao, G., Ma, S., Aafer, Y., Zhang, X.: Abs: scanning neural networks for back-doors by artificial brain stimulation. In: CCS 2019, pp. 1265–1282. ACM (2019)
19. Liu, Y., Ma, S., Aafer, Y., Lee, W.C., Zhai, J., Wang, W., Zhang, X.: Trojaning attack on neural networks (2017)
20. Liu, Y., Ma, X., Bailey, J., Lu, F.: Reflection backdoor: a natural backdoor attack on deep neural networks. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12355, pp. 182–199. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58607-2_11
21. Lu, J., Kumar, M.P.: Neural network branching for neural network verification. arXiv preprint arXiv:1912.01329 (2019)
22. Mirman, M., Gehr, T., Vechev, M.: Differentiable abstract interpretation for provably robust neural networks. In: ICML 2018, pp. 3578–3586. PMLR (2018)
23. Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. In: CVPR 2017, pp. 1765–1773. IEEE (2017)
24. Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A.: The limitations of deep learning in adversarial settings. In: EuroS&P 2016, pp. 372–387. IEEE (2016)

25. Parkhi, O.M., Vedaldi, A., Zisserman, A.: Deep face recognition (2015)
26. Pham, L.H., Li, J., Sun, J.: Socrates: Towards a unified platform for neural network verification. arXiv preprint arXiv:2007.11206 (2020)
27. Pham, L.H., Sun, J.: Source and Benchmark (2022). https://doi.org/10.6084/m9.figshare.19719742
28. Shafahi, A., Huang, W.R., Najibi, M., Suciu, O., Studer, C., Dumitras, T., Goldstein, T.: Poison frogs! targeted clean-label poisoning attacks on neural networks. arXiv preprint arXiv:1804.00792 (2018)
29. Singh, G., Ganvir, R., Püschel, M., Vechev, M.: Beyond the single neuron convex barrier for neural network certification. In: NeurIPS 2019, pp. 15098–15109 (2019)
30. Singh, G., Gehr, T., Mirman, M., Püschel, M., Vechev, M.: Fast and effective robustness certification. In: NeurIPS 2018, pp. 10802–10813 (2018)
31. Singh, G., Gehr, T., Püschel, M., Vechev, M.: Boosting robustness certification of neural networks. In: International Conference on Learning Representations (2018)
32. Singh, G., Gehr, T., Püschel, M., Vechev, M.: An abstract domain for certifying neural networks. Proceedings of the ACM on Programming Languages **3**, 1–30 (2019)
33. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199 (2013)
34. Tjeng, V., Xiao, K., Tedrake, R.: Evaluating robustness of neural networks with mixed integer programming. arXiv preprint arXiv:1711.07356 (2017)
35. Turner, A., Tsipras, D., Madry, A.: Clean-label backdoor attacks (2018)
36. Virtanen, P., et al.: SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. Nature Methods **17**, 261–272 (2020). https://doi.org/10.1038/s41592-019-0686-2
37. Wang, B., Cao, X., Gong, N.Z., et al.: On certifying robustness against backdoor attacks via randomized smoothing. arXiv preprint arXiv:2002.11750 (2020)
38. Wang, B., Yao, Y., Shan, S., Li, H., Viswanath, B., Zheng, H., Zhao, B.Y.: Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In: S&P 2019, pp. 707–723. IEEE (2019)
39. Wang, J., Dong, G., Sun, J., Wang, X., Zhang, P.: Adversarial sample detection for deep neural network through model mutation testing. In: ICSE 2019, pp. 1245–1256. IEEE/ACM (2019)
40. Wang, S., Pei, K., Whitehouse, J., Yang, J., Jana, S.: Efficient formal safety analysis of neural networks. In: NeurIPS 2018, pp. 6367–6377 (2018)
41. Wang, S., Pei, K., Whitehouse, J., Yang, J., Jana, S.: Formal security analysis of neural networks using symbolic intervals. In: 27th USENIX Security Symposium (USENIX Security 18), pp. 1599–1614 (2018)
42. Zhang, P., et al.: White-box fairness testing through adversarial sampling. In: ICSE 2020, pp. 949–960. IEEE/ACM (2020)