
Certified Interpretability Robustness for Class Activation Mapping

Alex Gu
MIT
Cambridge, MA, 02139
gua@mit.edu

Tsui-Wei Weng
MIT-IBM Watson AI Lab
IBM Research
tweng@mit.edu

Pin-Yu Chen
MIT-IBM Watson AI Lab
IBM Research
pin-yu.chen@ibm.com

Sijia Liu
MIT-IBM Watson AI Lab
IBM Research
sijia.liu@ibm.com

Luca Daniel
MIT
Cambridge, MA, 02139
luca@mit.edu

Abstract

Interpreting machine learning models is challenging but crucial for ensuring the safety of deep networks in autonomous driving systems. Due to the prevalence of deep learning based perception models in autonomous vehicles, accurately interpreting their predictions is crucial. While a variety of such methods have been proposed, most are shown to lack robustness. Yet, little has been done to provide certificates for interpretability robustness. Taking a step in this direction, we present CORGI, short for **C**ertifiably **p**rOvable **R**obustness **G**uarantees for **I**nterpretability mapping. CORGI is an algorithm that takes in an input image and gives a certifiable lower bound for the robustness of the top k pixels of its CAM interpretability map. We show the effectiveness of CORGI via a case study on traffic sign data, certifying lower bounds on the minimum adversarial perturbation not far from (4-5x) state-of-the-art attack methods.

1 Introduction

Deep neural networks are commonplace in safety-critical instances such as real-world autonomous driving systems. However, a major issue surrounding these networks is that they lack robustness to adversarial perturbations [19]. Under visually imperceptible perturbations, well-trained models can easily be fooled into confidently making predictions for an incorrect class, even in the black-box setting when the adversary does not even have access to the network architecture [8]. In the past few years, we have seen a series of increasingly stronger attacks against neural networks (Fast Gradient Sign Method [7], Projected Gradient Descent [10], Carlini-Wagner [3]).

A dual view of robustness research has focused on providing provable defenses against these attacks. For a neural network NN and some input image x , attacks aim to find an adversarial example x' such that $NN(x) \neq NN(x')$ and $\|x - x'\|_\infty = \epsilon$ is as small as possible. In contrast, the goal of these provable defenses is to provide a *certifiable radius* r that would guarantee that $NN(x) = NN(x')$ for all $\|x - x'\|_\infty \leq r$. While being able to determine the maximum certificate r^* would provide us optimal levels of safety, doing so, even under an approximation ratio, has been shown to be hard [21]. However, various approaches have been used to come close to this goal, such as semidefinite programming [12], mixed-integer programming [20], and linear bounding of activation functions [2].

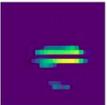
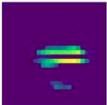
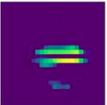
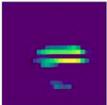
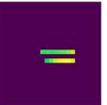
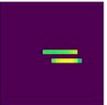
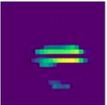
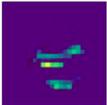
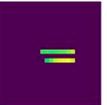
Attack Radius	Image		CAM		Top- k CAM Pixels		Pixels Same
	Original	Perturbed	Original	Perturbed	Original	Perturbed	
$\varepsilon = 0.0009$ (CORGI bound)							15 (all)
$\varepsilon = 0.002$							14
$\varepsilon = 0.07$							3

Figure 1: An image in the GTSRB dataset under 3 different attack radii. Shown are the original and perturbed images, CAM maps, and the top- k CAM pixel positions ($k = 15$). Under our CORGI certificate, all top- k pixels remain identical.

A frightening concern is that while deep neural networks make incorrect predictions on adversarial examples, they leave little to no insight towards what went wrong. As autonomous vehicle systems are comprised of deep neural networks for perception systems [4], being able to interpret their predictions is crucial in establishing trust and understanding between humans and these systems. This issue has given rise to many interpretability methods including CAM [25], GradCAM [13], Integrated Gradients [18], SmoothGrad [16], and DeepLIFT [14]. All of these methods generate an interpretability map for each image showing the effect of each input pixel on the overall prediction.

Unfortunately, recent work has shown that these interpretability methods are also prone to adversarial perturbations [24] [1] [6]. All three works successfully generate visually imperceptible perturbations for which the network correctly classifies the image, but shows a nonsensical interpretability map. While providing certifiable guarantees for adversarial attacks on a neural network classifiers has been thoroughly studied [15], [22], little has been done to certify robustness for interpretability methods.

The most relevant work to ours is [11], which certifies interpretability on a modification of SmoothGrad. While we follow their example and use the top- k metric for interpretability, our work focuses on CAM, a well-established visual interpretability algorithm that not only provides input attribution importance, but also visually localizes class-discriminative regions. CAM is better suited for autonomous vehicle applications, where being able to cleanly distinguish similar road signs is essential for accident prevention. For example, as shown in Figure 1, the CAM map shows that the neural network is seeing the white stripe in the no entry sign, which is unique to that class. In addition, while computation of SmoothGrad requires a backward pass for each test image, CAM only requires a forward pass. Given the importance of interpretability methods like CAM, we propose an algorithm for the interpretability robustness of CAM, as being able to guarantee a certain degree of safety is imperative for our trust in autonomous driving systems.

Our paper contains the following contributions:

- We present CORGI (Certifiably prOvable Robustness Guarantees for Interpretability mapping), the first algorithm giving a lower bound on the interpretability robustness for CAM.
- We provide two verifiable theoretical guarantees with corresponding experiments for an exact top- k metric and relaxed top- k metric, as well as a third theoretical result in terms of the Lipschitz constant of the neural network.
- We empirically compare our CORGI lower bounds to Ghorbani’s attack method’s upper bounds [6] on the GTSRB traffic [17] dataset, showing that on average, our CORGI lower bounds are only a factor of 4-5x away from current attack upper bounds.

2 Background and Notation

In this section, we introduce our notation, explain how CAM interpretability maps are generated, and describe the attack algorithm [6] we use as an complementary upper bound.

2.1 Notation

For a neural network, input image $x \in \mathbb{R}^{d \times d}$, and arbitrary class c , let $I^c(x) \in \mathbb{R}^{s \times s}$ be the *interpretability map* for input image x and class c generated by CAM (described in Section 2.2). In this work, we are interested in the CAM map only for the predicted class, so we will often drop the c and write $I(x)$, understood to mean the interpretability map of input image x for its predicted class.

In addition, let $I(x)_1, I(x)_2, \dots, I(x)_{s^2} \in \mathbb{R}$ be the values of the individual pixels of $I(x)$, and $\sigma_1, \sigma_2, \dots, \sigma_{s^2}$ be a permutation of $1, 2, \dots, s^2$ such that $I(x)_{\sigma_1} \geq I(x)_{\sigma_2} \geq \dots \geq I(x)_{\sigma_{s^2}}$. Therefore, the σ_i 's denotes indices of the pixels in sorted order. For ease of notation, we will use $I(x)_{\sigma_i}$ interchangeably with $I(x)_{[i]}$ to denote the value of the i th largest pixel of the interpretability map, with σ_i being the position of that pixel in the original map.

2.2 Class Activation Mapping

One popular technique used to generate interpretability maps for deep neural networks is known as *class activation mapping* [25]. A class activation map (CAM) is a 2-dimensional heatmap that explains a neural network's prediction on an image by highlighting the parts of that image contributing to its prediction. A major advantage of CAM is that they can visually localize class-discriminative regions, which helps in understanding how a network distinguishes between similar classes.

CAM only provides an interpretability map for convolutional neural network architectures with their last convolutional layer followed by a global average pooling layer and a fully connected layer. Formally, consider such a network with last convolutional layer output of size $s \times s \times f$, where $s \times s$ is the size of the map, and f is the number of filters. Fix an input image, and let $f_{x,y}^k \in \mathbb{R}$ represent the value of spatial location (x, y) of filter k in the last convolutional layer. Let $F^k \in \mathbb{R}$ be the output of the k th filter after the global average pooling layer,

$$F^k = \sum_{1 \leq x, y \leq s} f_{x,y}^k$$

Now, let $w \in \mathbb{R}^k$ be the weights of the fully connected layer connected to the predicted class, which we denote as c . Therefore, the input to the final softmax is $\sum_k w_k^c F^k$, where w_k^c are weights that intuitively capture the importance of feature map k for the predicted class c . As a notational note, we will sometimes write this as w_k when it is clear that c is the predicted class. The class activation map for class c , $M^c \in \mathbb{R}^{s \times s}$, is computed as a weighted average of these feature maps:

$$M_{x,y}^c = \sum_k w_k^c f_{x,y}^k$$

2.3 Iterative Feature Importance Attack

One property that we would like interpretability maps to have is that the top- k largest pixels still remain fairly large under small perturbations. Denoting by σ_i the position of the i th largest pixel in $I(x)$ and σ'_i likewise in $I(x + \epsilon)$, we would like $\{\sigma_1, \sigma_2, \dots, \sigma_k\} = \{\sigma'_1, \sigma'_2, \dots, \sigma'_k\}$. In other words, we want the k largest pixels of the original image to remain the k largest pixels in the perturbed image. A good attack would then ensure the number of σ_i 's that appear in $\{\sigma'_1, \sigma'_2, \dots, \sigma'_k\}$ is small.

The iterative feature importance attack [6] intuitively attacks features maps under the top- k metric by finding a perturbation minimizing the sum of the values of the original top- k pixel positions in the new image. Given an input image x , their attack first considers $B = \{\sigma_1, \sigma_2, \dots, \sigma_k\}$, the top- k pixel locations in $I(x)$. They then run projected FGSM (fast-gradient sign method) [7] to minimize the objective $\sum_{i \in B} I(y)_i$ or equivalently, to maximize $D(x, y) = -\sum_{i \in B} I(y)_i$. In ReLU networks, however, where $\nabla_y D(x, y_{i-1})$ is piecewise-constant, we first replace the ReLU with a softplus as the original paper does. Therefore, $I(\cdot)$ is actually the interpretability map of the softplus network rather than the original ReLU network.

However, we make a small modification of their algorithm: at the end, instead of returning the y with the largest dissimilarity, we return the y such that $I(x)$ and $I(y)$ have the least number of top- k pixels in common, as this is what CORGI certifies for. The algorithm is as follows:

Algorithm 1 Gradient-based top- k interpretability map attack [6]

```

1: procedure TOP-K ATTACK( $x, \epsilon, I(\cdot), k$ )
2:   Define  $B = \{\sigma_1, \sigma_2, \dots, \sigma_k\}$  to be the top- $k$  pixel locations of  $I(x)$ 
3:   Define  $D(x, y) = -\sum_{i \in B} I(y)_i$  (dissimilarity metric between interpretability maps of  $x, y$ )
4:   Initialize  $y_0 = x$ 
5:   for  $i = 1, \dots, P$  do
6:      $y_i = y_{i-1} + \alpha \cdot \text{sign}(\nabla_y D(x, y_{i-1}))$ 
7:     If necessary, clip  $y_i$  to satisfy  $\|y_i - x\|_\infty \leq \epsilon$ 
8:   end for
9:   return  $y \in \{y_0, \dots, y_P\}$  such that  $I(x)$  and  $I(y)$  have the least number of top- $k$  pixels in
      common
10: end procedure

```

3 CORGI: A Framework for Certifiable Interpretability Robustness

In this section, we present CORGI, a framework which provides a lower bound r for the minimum adversarial attack radius r^* . When the interpretability attack method described in Section 2.3 is successful and returns an adversarial example $\|x' - x\|_\infty = \epsilon$, ϵ is then an upper bound for r^* . In contrast, we give a lower bound certificate r such that for perturbed images x' with $\|x - x'\|_\infty \leq r$, the set of top- k pixel positions of the original CAM map does not change. In other words, if $\{\sigma'_i\}$ denotes the top k pixel positions of the CAM map $I(x')$, then $\{\sigma_1, \dots, \sigma_k\} = \{\sigma'_1, \dots, \sigma'_k\}$. Note that the ordering of these pixel positions does not necessarily need to be the same, just that the unordered set stay constant. One subtlety to mention is that our CORGI bound does not guarantee that the classifier itself is robust. However, we can take the minimum of our CORGI bound and the robustness bound to give a certificate that classification and interpretation are both accurate. Empirically, however, robust interpretability seems to be a more challenging task than robust classification, and we observed that CORGI bounds are almost always smaller than robustness bounds.

CORGI builds off of the CNN-Cert framework [2], which provides efficient layer-wise bounds on convolutional neural networks. At a high level, CORGI operates in 2 stages. First, it uses CNN-Cert to give lower and upper bounds $L(\delta), R(\delta) \in \mathbb{R}^{s \times s}$ on the CAM map:

$$L(\delta)_{i,j} \triangleq \min_{\|x' - x\|_\infty \leq \delta} I(x')_{i,j}, U(\delta)_{i,j} \triangleq \max_{\|x' - x\|_\infty \leq \delta} I(x')_{i,j}$$

Now, observe that guaranteed top- k interpretability robustness under perturbations of radius at most r implies robustness under any smaller perturbation radius $r' \leq r$. Therefore, in the second stage, CORGI uses binary search to find the largest radius r that it is able to accurately certify, giving us a lower bound certificate on interpretability robustness.

3.1 Upper and Lower Bounding CAM Values with CNN-Cert

Given a neural network and an input, CNN-Cert is a framework that is able to lower and upper bound all of the layers of a neural network under a maximum perturbation radius δ . We use CNN-Cert as a subroutine to lower and upper bound CAM maps. First, we lower and upper bound the last convolutional layer output. Let $l_k(\delta), u_k(\delta) \in \mathbb{R}^{s \times s}$ denote the bounds of the k th filter of the last convolutional layer output f^k , so that $l_k(\delta)_{i,j} \leq f^k(x')_{i,j} \leq u_k(\delta)_{i,j}$ for all pixels (i, j) and perturbed images x' with $\|x' - x\|_\infty \leq \delta$. We then combine these lower/upper bounds with the known weights w_k of the fully connected layer to bound the CAM. To see how to do this, observe

that for all $\|x' - x\|_\infty \leq \delta$,

$$\begin{aligned}
I(x')_{i,j} &= \sum_k w_k f^k(x')_{i,j} \\
&= \sum_{k, w_k < 0} w_k f^k(x')_{i,j} + \sum_{k, w_k > 0} w_k f^k(x')_{i,j} \\
&\geq \sum_{k, w_k < 0} w_k u_k(\delta)_{i,j} + \sum_{k, w_k > 0} w_k l_k(\delta)_{i,j} \triangleq L(\delta)_{i,j}
\end{aligned} \tag{1}$$

Similarly, we have

$$I(x')_{i,j} \leq \sum_{k, w_k < 0} w_k l_k(\delta)_{i,j} + \sum_{k, w_k > 0} w_k u_k(\delta)_{i,j} \triangleq U(\delta)_{i,j} \tag{2}$$

Therefore, with these definitions of $L(\delta), U(\delta) \in \mathbb{R}^{s \times s}$, we then have, for all $\|x' - x\|_\infty \leq \delta$, we now have an upper and lower bound on each pixel of the CAM map:

$$L(\delta)_{i,j} \leq I(x')_{i,j} \leq U(\delta)_{i,j} \tag{3}$$

3.2 Testing for Certifiability

The next stage of CORGI aims to find the largest perturbation radius that can be certified with the bounds in Eq. 3. For a fixed input image x and a perturbation size δ , a perturbed image x' with $\|x - x'\|_\infty \leq \delta$ will have the same top- k pixels if and only if

$$\min_{k' \in 1, 2, \dots, k} I(x')_{\sigma_{k'}} \geq \max_{k' \in k+1, \dots, s^2} I(x')_{\sigma_{k'}} \tag{4}$$

We will find a condition that ensures that Eq. 4 holds. First, observe that since $I(x')_{\sigma_k} \geq L(\delta)_{\sigma_k}$ for all σ_k ,

$$\min_{k' \in 1, 2, \dots, k} I(x')_{\sigma_{k'}} \geq \min_{k' \in 1, 2, \dots, k} L(\delta)_{\sigma_{k'}} \tag{5}$$

Similarly, for the upper bound,

$$\max_{k' \in k+1, \dots, s^2} U(\delta)_{\sigma_{k'}} \geq \max_{k' \in k+1, \dots, s^2} I(x')_{\sigma_{k'}} \tag{6}$$

Eq. 5 and 6 then give rise to our theorem for certified CAM interpretability:

Theorem 3.2 (CAM Interpretability Certificate) *Let $x \in \mathbb{R}^{s \times s}$ be a fixed input image, $\delta \in \mathbb{R}_{\geq 0}$ be a fixed maximum perturbation size, $k \in \mathbb{Z}_{\geq 1}$, and $L(\delta), R(\delta)$ be CAM map bounds as computed in Eq. 1. Then, for all x' such that $\|x - x'\|_\infty \leq \delta$, if*

$$\min_{k' \in 1, 2, \dots, k} L(\delta)_{\sigma_{k'}} \geq \max_{k' \in k+1, \dots, s^2} U(\delta)_{\sigma_{k'}} \tag{7}$$

then δ is a certified lower bound of r^ . In other words, the top- k pixel positions of the CAM map $I(x')$ must be the same as those in the original CAM map $I(x)$.*

Proof This theorem follows since Eq. 4 follows directly from Eq. 5, 6, and the hypothesis Eq. 7:

$$\min_{k' \in 1, 2, \dots, k} I(x')_{\sigma_{k'}} \geq \min_{k' \in 1, 2, \dots, k} L(\delta)_{\sigma_{k'}} \geq \max_{k' \in k+1, \dots, s^2} U(\delta)_{\sigma_{k'}} \geq \max_{k' \in k+1, \dots, s^2} I(x')_{\sigma_{k'}} \tag{8}$$

As a remark, observe that for $\delta = 0$, the inequality is obviously true because CNN-Cert will return the values of the network at the original input, so $L(0) = I(x) = U(0)$, and $I(x)_{[k]} = \min_{k' \in 1, 2, \dots, k} L(0)_{\sigma_{k'}} \geq \max_{k' \in k+1, \dots, s^2} U(0)_{\sigma_{k'}} = I(x)_{[k+1]}$ by definition.

Now, notice that since the lower and upper bounds of CNN-Cert are monotonic in δ , if $\delta' \leq \delta$, then $L(\delta)_{i,j} \leq L(\delta')_{i,j}$ and $U(\delta')_{i,j} \leq U(\delta)_{i,j}$, so that $L(\delta)_{\sigma_i} \leq L(\delta')_{\sigma_i}$ and $U(\delta')_{\sigma_i} \leq U(\delta)_{\sigma_i}$ for all $1 \leq i \leq s^2$. Therefore, if Eq. 7 is satisfied for some δ , then it is also satisfied for all $\delta' \geq \delta$. We can then binary search for the smallest certifiable radius.

Algorithm 2 CORGI

```
1: procedure CORGI( $x, f(\cdot), k$ )
2:   Use  $f$  to compute the interpretability map  $I(x)$ 
3:   Let  $\sigma_1, \sigma_2, \dots, \sigma_{s^2}$  be the indices of  $I(x)$ 's pixels from largest to smallest
4:   Initialize  $\delta_l \leftarrow 0, \delta_u \leftarrow 1$ 
5:   for  $i = 1, \dots, P$  do
6:      $\delta_m \leftarrow (\delta_l + \delta_u)/2$ 
7:     Run CNN-Cert( $x, \delta_m$ ) to get  $l_k(\delta_m), u_k(\delta_m)$ , last convolutional output layer bounds
8:     Compute  $L(\delta_m), U(\delta_m)$  from  $l_k(\delta_m), u_k(\delta_m)$ , and  $f$ 's weights  $w_k$  (Eq. 1, 2)
9:     if Eq. 7 holds (which means  $\delta_m$  is a certificate) then
10:        $\delta_u \leftarrow \delta_m$ 
11:     else
12:        $\delta_l \leftarrow \delta_m$ 
13:     end if
14:   end for
15:   return  $\delta_u$  (an lower bound on the certificate)
16: end procedure
```

3.3 A More Relaxed Interpretability Guarantee

While requiring that the top- k pixel positions remain the same is a nice condition to have, sometimes it will be enough for the top- k pixel positions to remain relatively large in the perturbed image without remaining in the top- k . If $I(x)_{[k]}$ is very close in value to $I(x)_{[k+1]}$, requiring that the original top- k pixels $\{\sigma_1, \dots, \sigma_k\}$ be the same as the perturbed top- k pixels $\{\sigma'_1, \dots, \sigma'_k\}$ may be a strong condition. In this case, even if a perturbation swapped the relative values of pixels σ_k and σ_{k+1} causing $I(x + \delta)_{\sigma_k} < I(x + \delta)_{\sigma_{k+1}}$, the interpretability map would still remain fairly accurate. Therefore, we relax the original condition and instead only require that $\{\sigma_1, \dots, \sigma_k\} \subseteq \{\sigma'_1, \dots, \sigma'_{k_2}\}$, for some $k_2 \geq k$. Intuitively, this says that the top- k pixel positions of the original image must remain in the top- k_2 pixel positions of the perturbed image. Note that $k_2 = k$ corresponds to our original requirement that the top- k pixel positions remain the same under perturbation.

The algorithm for this scenario remains very similar to that in Section 3.2, and the crux is once again checking whether or not we can certify top- k interpretability robustness for a given δ . As before, we first compute $L(\delta), U(\delta)$ as in Section 3.1. Now, define $\tau_1, \dots, \tau_{s^2-k}$ to be an ordering of all pixel positions in $S = \{1, 2, \dots, s^2\} \setminus \{\sigma_1, \sigma_2, \dots, \sigma_k\}$ satisfying $U(\delta)_{\tau_1} \geq U(\delta)_{\tau_2} \geq \dots \geq U(\delta)_{\tau_{s^2-k}}$, and for notational simplicity define $U(\delta)_{\{k\}} = U(\delta)_{\tau_k}$. $U(\delta)_{\{k\}}$ then denotes the k th largest upper bound of all the pixels that weren't originally in the top- k of $I(x)$. Similarly, define $U(\delta)_{[k]}$ to be the k th largest pixel value in $U(\delta)$ (no longer restricted on S) and $I(x')_{[k]}$ to be the k th largest pixel value in the interpretability map $I(x')$. We can then prove a theorem analogous to Theorem 3.2:

Theorem 3.3 (Relaxed CAM Interpretability Certificate) *Let $x \in \mathbb{R}^{s \times s}$ be a fixed input image, $\delta \in \mathbb{R}_{\geq 0}$ be a fixed maximum perturbation size, $k, k_2 \in \mathbb{Z}_{\geq 1}$ such that $k_2 \geq k$, and $L(\delta), R(\delta)$ be CAM map bounds as computed in Eq. 1. Then, for all x' such that $\|x - x'\|_{\infty} \leq \delta$, if*

$$\min_{k' \in \{1, 2, \dots, k\}} L(\delta)_{\sigma_{k'}} \geq U(\delta)_{\{k_2 - k + 1\}} \quad (9)$$

then under all perturbations of size at most δ , the top- k pixel positions of the original CAM map $I(x)$ are in the top k_2 pixel positions of the CAM map $I(x')$.

The proof is very similar to that of Theorem 3.2 and is deferred to the appendix. Using this theorem, we can simply run the same binary search as described in the previous section. The only modification that must be made to the original algorithm is that we must now check Eq. 9 instead of Eq. 7 in line 8 of CORGI.

4 Experiments

Dataset and Model We use the German Traffic Sign Recognition Benchmark (GTSRB) [17], a large, lifelike database of over 50000 $48 \times 48 \times 3$ traffic signs in 43 classes. We train a small

convolutional neural network of four convolutional layers with 8, 8, 16, and 16 filters, a global average pooling layer, and a fully connected layer (as required by CAM).

Interpretability Attack First, we ran CORGI with $k = 15$ on a sample "no entry" sign to get a certified radius of $r = 0.0009$. We first implemented the attack algorithm (Algorithm 1) with $P = 300$ and $\alpha = \epsilon/10$. Figure 1 shows the attack algorithm’s perturbations on both r and two larger values of ϵ . The perturbed images, CAM maps, and top- k maps are shown. Note that within our CORGI bound, all the top- k pixels remain the same.

Comparing CORGI Bounds to Interpretability Attack Bounds By running binary search on the interpretability attack algorithm, we can find the maximum ϵ that yields a successful top- k attack, giving us an upper bound on the best certificate r^* . Since CORGI is currently the first framework that provides interpretability robustness certificates for CAM, we compare CORGI’s certificates with complementary upper bounds generated by the iterative feature importance attack. To do so, we sampled 100 random images from each of two classes, class 17 ("no entry") and class 12 ("priority road"). Figure 2a shows two samples of each image.

Figure 2b shows both the CORGI lower bound and attack upper bound for each of the 200 images. For ease of comparison, the images in each class were sorted by their CORGI lower bound. For each of the 200 images, we also calculated the ratio ϵ/r , where ϵ is the minimum perturbation radius found by the attack, and r is the maximum CORGI certifiable radius (Figure 2c). There were a few outliers in both classes, where the attack bound was abnormally large. The figure removes these outliers and only plots samples for which the gap was under 20. The resulting figure consists of 99/100 class 12 samples and 92/100 of the class 17 samples. the majority of the ratios were fairly small, with a median of 5.0 for class 17 and 4.7 for class 12. This signifies that for this network, CORGI bounds are only 4-5 times larger than attack bounds, and therefore from the optimal perturbation r^* . We show additional examples of CORGI vs attack bounds for 5 other classes in the appendix.

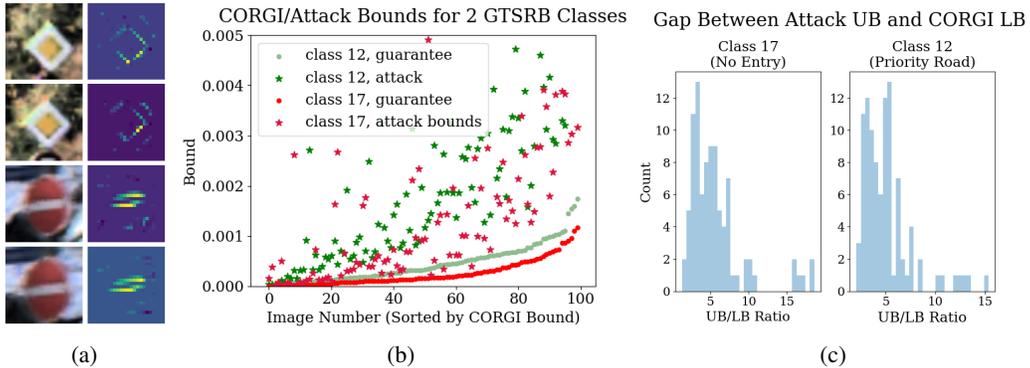


Figure 2: (a) Sample images and CAM maps from class 12 (priority road) and class 17 (no entry). (b) Comparison of CORGI lower bounds and attack bounds for 100 random samples for each class. (c) Histogram of the ratio of attack upper bounds to CORGI lower bounds, with a median of 4.7 for class 12 and 5.0 for class 17.

Effectively choosing values of k for interpretability maps As people seek to produce accurate and descriptive interpretability maps, one important question is how k should be chosen to faithfully represent the important regions of an input image. If k is too small, then the top- k CAM map may not fully capture the salient regions. If k is too large, then we might capture noise pixels that were not pertinent to the model’s prediction. Figure 3(a) shows a sample image and its top- k CAM map with increasing k values.

One natural way to evaluate a specific value of k is to consider $I(x)_{[k]} - I(x)_{[k+1]}$, which we call the *CAM gap* for position k . Intuitively, if for some k , $I(x)_{[k]}$ and $I(x)_{[k+1]}$ are close, then only a small perturbation should be needed for $I(x)_{[k]}$ and $I(x)_{[k+1]}$ to swap places. Since the k th and $(k + 1)$ st largest pixel might not be well distinguished enough, that value of k might be a poor choice. We would also expect that such a k would be easily attacked, and more generally, that top- k interpretability maps for values of k with a larger CAM gap will require a larger perturbation to successfully attack.

We varied k from 1 to 50 in order to see how the CAM map gap was correlated to the performance of both our algorithm and the attack. Figure 3b demonstrates that as the CAM gap increases, the CORGI certificates and attack bounds both increase. This captures the intuition that the larger the CAM gap for a value k , the more robust the top- k CAM interpretability map will be.

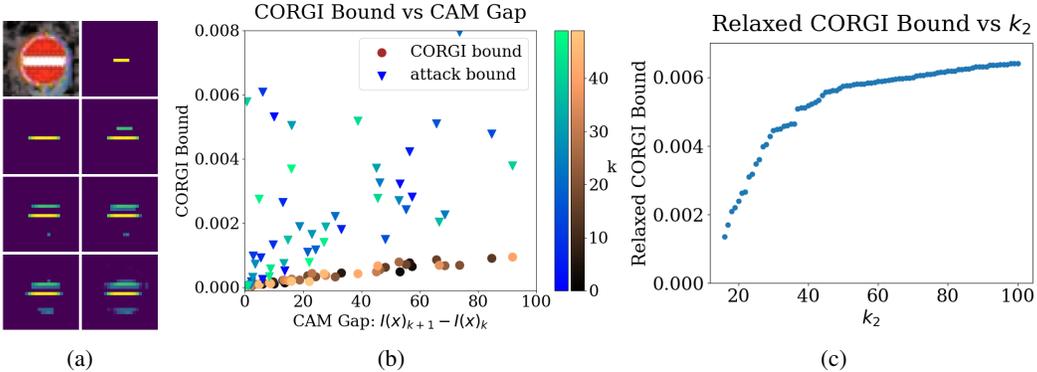


Figure 3: (a) The top- k CAM map for an image with increasing values of k from 5 to 70. (b) Both CORGI certificates and attack bounds are highly correlated with the CAM gap, meaning that choosing a k with a large CAM gap leads to better interpretability robustness. (c) Relaxed CORGI bounds as a function of k_2 , the maximum position of an original top- k pixel in the perturbed image.

Relaxing the top- k condition We ran the modified CORGI algorithm given in Section 3.3 for a fixed image and fixed k . We varied the threshold k_2 to see how increasing k_2 would affect CORGI’s bound. The results are shown in Figure 3c. For $k_2 \approx k$, increasing k_2 by just a small amount is able to strengthen our certificate by quite a bit. However, as k_2 becomes larger, the CORGI bounds begin to level off. We believe that this occurs because

5 Conclusions and Future Work

We have presented CORGI, the first certifiable interpretability robustness for CAM. By using CNN-Cert to compute last layer convolutional output bounds, CORGI is able to efficiently provide a certificate $r \leq r^*$. Through comparing our bounds with upper bounds generated by the iterative feature importance attack algorithm, CORGI provides bounds that are not far from certifiable upper bounds and acts as a solid baseline for future work in interpretability robustness.

A limitation of CORGI is that it only performs well for small networks. Even for small networks, CORGI takes about 1 minute per image. GTSRB images are significantly lower in resolution than images seen by autonomous vehicles, and networks used in the area are generally much larger, so CORGI does not yet scale to real-world applications. The main accuracy bottleneck of CORGI lies in the absence of fast and tight bounds for large convolutional neural networks. Therefore, improving both layer-wide network bounds and interpretability bounds for larger networks is a crucial direction to explore.

We have not yet addressed certifiable robustness for other interpretability methods past CAM, such as DeepLIFT, Integrated Gradients, and GradCAM. Since CAM is a bit restrictive in the architectures that it is able to handle and requires a global average pooling layer following convolutional layers, developing bounds for these other interpretability algorithms will expand the class of models for which we can develop certificates.

6 Acknowledgement

Alex Gu, Tsui-Wei Weng and Luca Daniel are partially supported by MIT-IBM Watson AI Lab.

7 Bibliography

- [1] David Alvarez-Melis and Tommi S Jaakkola. On the robustness of interpretability methods. *arXiv preprint arXiv:1806.08049*, 2018.
- [2] Akhilan Boopathy, Tsui-Wei Weng, Pin-Yu Chen, Sijia Liu, and Luca Daniel. Cnn-cert: An efficient framework for certifying robustness of convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3240–3247, 2019.
- [3] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [4] Jianyu Chen, Zhuo Xu, and Masayoshi Tomizuka. End-to-end autonomous driving perception with sequential latent representation learning. *arXiv preprint arXiv:2003.12464*, 2020.
- [5] Mahyar Fazlyab, Alexander Robey, Hamed Hassani, Manfred Morari, and George Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, pages 11423–11434, 2019.
- [6] Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3681–3688, 2019.
- [7] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [8] Chuan Guo, Jacob R Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Q Weinberger. Simple black-box adversarial attacks. *arXiv preprint arXiv:1905.07121*, 2019.
- [9] Matt Jordan and Alexandros G Dimakis. Exactly computing the local lipschitz constant of relu networks. *arXiv preprint arXiv:2003.01219*, 2020.
- [10] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016.
- [11] Alexander Levine, Sahil Singla, and Soheil Feizi. Certifiably robust interpretation in deep learning. *arXiv preprint arXiv:1905.12105*, 2019.
- [12] Aditi Raghunathan, Jacob Steinhardt, and Percy S Liang. Semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in Neural Information Processing Systems*, pages 10877–10887, 2018.
- [13] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [14] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. *arXiv preprint arXiv:1704.02685*, 2017.
- [15] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–30, 2019.
- [16] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [17] Johannes Stalldkamp, Marc Schlipf, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012.
- [18] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. *arXiv preprint arXiv:1703.01365*, 2017.

- [19] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [20] Vincent Tjeng, Kai Xiao, and Russ Tedrake. Evaluating robustness of neural networks with mixed integer programming. *arXiv preprint arXiv:1711.07356*, 2017.
- [21] Tsui-Wei Weng, Huan Zhang, Hongge Chen, Zhao Song, Cho-Jui Hsieh, Duane Boning, Inderjit S Dhillon, and Luca Daniel. Towards fast computation of certified robustness for relu networks. *arXiv preprint arXiv:1804.09699*, 2018.
- [22] Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. In *Advances in neural information processing systems*, pages 4939–4948, 2018.
- [23] Huan Zhang, Pengchuan Zhang, and Cho-Jui Hsieh. Recurjac: An efficient recursive algorithm for bounding jacobian matrix of neural networks and its applications. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5757–5764, 2019.
- [24] Haizhong Zheng, Earleence Fernandes, and Atul Prakash. Analyzing the interpretability robustness of self-explaining models. *arXiv preprint arXiv:1905.12429*, 2019.
- [25] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.

8 Appendix

8.1 Model and Experimental Details

The architecture of our neural network is shown below. Note that the CAM map is a weighted average of the 16 filters in the last convolutional layer, so has shape 24×24 . The test accuracy of our network is around 92%. Ghorbani’s attack algorithm was run with $P = 300$ steps and $\alpha = \epsilon/10$. Increasing P did not have a significant impact on the effectiveness of the attack.

Input ($48 \times 48 \times 3$)
Conv with 8 filters ($48 \times 48 \times 8$)
ReLU
Conv with 8 filters ($48 \times 48 \times 8$)
ReLU
Max Pool ($24 \times 24 \times 8$)
Dropout (0.2)
Conv with 16 filters ($24 \times 24 \times 16$)
ReLU
Conv with 16 filters ($24 \times 24 \times 16$)
ReLU
Global Average Pooling (16)
Dense (43)

8.2 Relaxed CAM Certificate

In this section, we give the proof of Theorems 3.3 in our main work.

Theorem 3.3 (Relaxed CAM Interpretability Certificate) *Let $x \in \mathbb{R}^{s \times s}$ be a fixed input image, $\delta \in \mathbb{R}_{\geq 0}$ be a fixed maximum perturbation size, $k, k_2 \in \mathbb{Z}_{\geq 1}$ such that $k_2 \geq k$, and $L(\delta), R(\delta)$ be CAM map bounds as computed in Eq. 1. Then, for all x' such that $\|x - x'\|_\infty \leq \delta$, if*

$$\min_{k' \in \{1, 2, \dots, k\}} L(\delta)_{\sigma_{k'}} \geq U(\delta)_{\{k_2 - k + 1\}} \quad (10)$$

then under all perturbations of size at most δ , the top- k pixel positions of the original CAM map $I(x)$ are in the top k_2 pixel positions of the CAM map $I(x')$.

Proof The proof is very similar to that of Theorem 3.2. We have the following chain of inequalities:

$$\min_{k' \in \{1, 2, \dots, k\}} I(x')_{\sigma_{k'}} \stackrel{(1)}{\geq} \min_{k' \in \{1, 2, \dots, k\}} L(\delta)_{\sigma_{k'}} \stackrel{(2)}{\geq} U(\delta)_{\{k_2 - k + 1\}} \stackrel{(3)}{=} U(\delta)_{[k_2 + 1]} \stackrel{(4)}{\geq} I(x')_{[k_2 + 1]}$$

(1) comes from Eq. 3, (2) comes from the hypothesis, and (4) is an implication of Eq. 1. To see why (3) holds, note that the hypothesis guarantees that $U(\delta)_{\sigma_1}, \dots, U(\delta)_{\sigma_k} \geq U(\delta)_{\{k_2 - k + 1\}}$, which means that the $(k_2 + 1)$ st largest pixel value of $U(\delta)$ must be equal to the $(k_2 - k + 1)$ st largest value of $U(\delta)$ without considering the σ_i ’s.

8.3 Lipschitz Constant CAM Certificate

In this section, we give a different theoretical bound directly based on the Lipschitz constant of a neural network.

First, recall the definition of a locally Lipschitz function:

Definition 3.4.1: A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -locally Lipschitz in a region $S \subseteq \mathbb{R}^d$ if for all $x, y \in S$, we have $|f(x) - f(y)|_1 \leq L\|x - y\|_\infty$. The smallest L for which this inequality is true is called the *local Lipschitz constant* of f in a region S .

We can then guarantee the following certificate, the proof of which will be deferred to the Appendix.

Theorem 8.3 (Lipschitz-Based CAM Interpretability Certificate) *Let $x \in \mathbb{R}^{s \times s}$ be a fixed input image, $I(\cdot) : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{s \times s}$ be the CAM interpretability mapping function, $k \in \mathbb{Z}_{\geq 1}$, and $\{\sigma_i\}_{1 \leq i \leq k}$ denote the top- k pixels of the input image $I(x)$. For $1 \leq i \leq k$, $k+1 \leq p \leq s^2$, define $g_p^{(i)}(x) \triangleq I(x)_{\sigma_i} - I(x)_p$, and let $L_p^{(i)}(\delta)$ be the local Lipschitz constant of the function $g_p^{(i)}$ in some region $S = \{x' | \delta \geq \|x - x'\|_\infty\}$ for some $\delta \in \mathbb{R}$. Finally, let*

$$r = \min_{k+1 \leq p \leq s^2} \left(\min_{1 \leq i \leq k} \frac{g_p^{(i)}(x)}{L_p^{(i)}(\delta)} \right)$$

Then, for all x such that

$$\|x - x'\|_\infty \leq r$$

the set of top- k pixels of $I(x)$ will be identical to that of $I(x')$ as long as δ is chosen such that $\delta \geq r$. In other words, r is a certified radius for top- k interpretability.

Proof Note that in order to have the top- k pixel positions of $I(x)$ be identical to that of $I(x')$, we must have, for all $1 \leq i \leq k$ and $k+1 \leq p \leq s^2$,

$$g_p^{(i)}(x') = I(x')_{\sigma_i} - I(x')_p \geq 0$$

This is essentially a restatement of Eq. 4. By the definition of the Lipschitz constant, for any ϵ such that $\|\epsilon\|_\infty < \delta$, we have

$$\begin{aligned} |g_p^{(i)}(x + \epsilon) - g_p^{(i)}(x)|_1 &\leq L_p^{(i)}(\delta) \|\epsilon\|_\infty \\ \Rightarrow g_p^{(i)}(x) - L_p^{(i)}(\delta) \|\epsilon\|_\infty &\leq g_p^{(i)}(x + \epsilon) \end{aligned}$$

Observe that if $\|\epsilon\|_\infty \leq \frac{g_p^{(i)}(x)}{L_p^{(i)}(\delta)}$, then $g_p^{(i)}(x + \epsilon) \geq 0$. Therefore, if

$$\|\epsilon\|_\infty \leq \min_{k+1 \leq p \leq s^2} \left(\min_{1 \leq i \leq k} \frac{g_p^{(i)}(x)}{L_p^{(i)}(\delta)} \right)$$

then $g_p^{(i)}(x + \epsilon) \geq 0$ for all $1 \leq i \leq k$, $k+1 \leq p \leq s^2$, which is precisely the condition that the top- k pixels of $I(x)$ are identical to the top- k pixels of $I(x + \epsilon)$ under $\|\epsilon\|_\infty \leq r$.

In its essence, CORGI provides certificates by relying on CNN-Cert to derive bounds on the interpretability maps and then ensuring that the lower bound of $I(x')_{\sigma_i}$ is larger than the upper bound of the other pixels. For larger networks like AlexNet and ResNets, however, CNN-Cert’s bounds are often from known attack bounds [2], and are suspected to be loose. Therefore, Theorem 8.3 provides a more straightforward approach by directly operating on the Lipschitz constant of the function $g_p^{(i)}$. Therefore, for tight estimates of the Lipschitz constant, we expect this bound to outperform CORGI. However, Lipschitz constant estimation is currently still a nascent and young area of research [9] [23] [5]. Therefore, as current methods for estimating Lipschitz constants are either very loose or run very slowly, we will not give an empirical analysis of this approach.

8.4 Additional Image Examples

We show two more examples of images, CORGI bounds, and CAM maps similar to Figure 1.

8.5 Additional CORGI Examples

We show a comparison between attack and CORGI bounds for 5 additional classes, similar to that of Figure 2. For each of the 5 classes, 100 test images were random chosen and evaluated. The distribution of ratios is similar for all of the classes, suggesting that CORGI performs fairly uniformly throughout the entire testing distribution.

Attack Radius	Image		CAM		Top- k CAM Pixels		Pixels Same
	Original	Perturbed	Original	Perturbed	Original	Perturbed	
$\epsilon = 0.0003$ (CORGI bound)							13/13
$\epsilon = 0.003$							12/13
$\epsilon = 0.07$							3/13

Attack Radius	Image		CAM		Top- k CAM Pixels		Pixels Same
	Original	Perturbed	Original	Perturbed	Original	Perturbed	
$\epsilon = 0.001$ (CORGI bound)							10/10
$\epsilon = 0.008$							9/10
$\epsilon = 0.07$							3/10

Figure 4: GTSRB images under 3 different attack radii. Shown are the original and perturbed images, CAM maps, and the top- k CAM pixel positions.

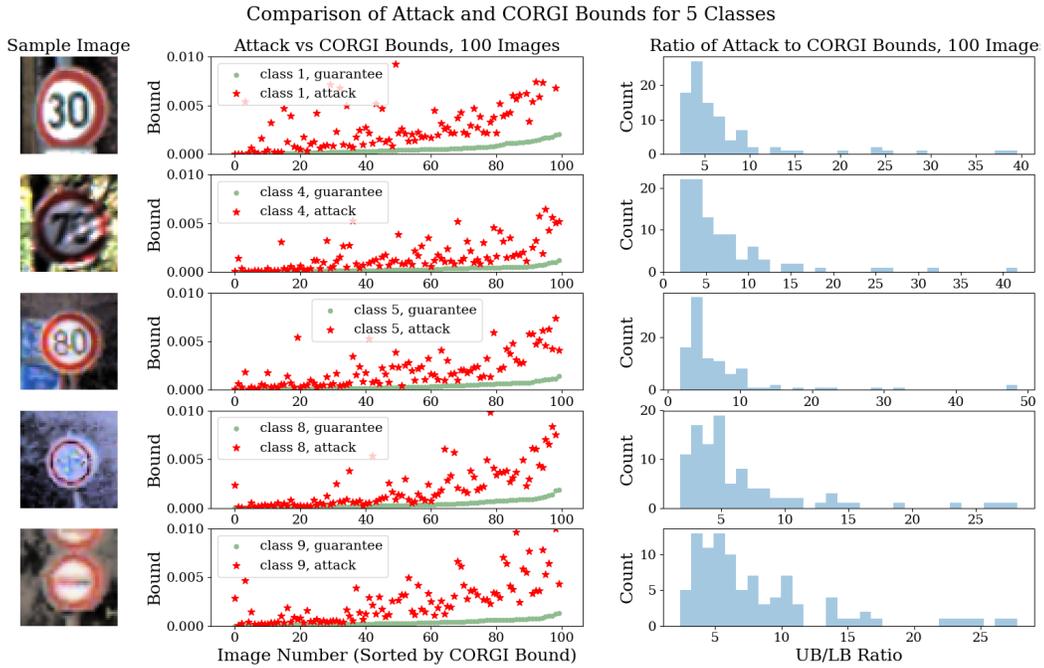


Figure 5: Comparison of CORGI vs attack bounds for 5 additional classes.