

Incremental Bounded Model Checking of Artificial Neural Networks in CUDA

Luiz H. Sena¹, Iury V. Bessa¹, Mikhail R. Gadelha², Lucas C. Cordeiro³, and Edjard Mota¹

¹Federal University of Amazonas, Manaus, Brazil

²SIDIA Instituto de Ciência e Tecnologia, Manaus, Brazil

³University of Manchester, Manchester, United Kingdom

Abstract—Artificial Neural networks (ANNs) are powerful computing systems employed for various applications due to their versatility to generalize and to respond to unexpected inputs/patterns. However, implementations of ANNs for safety-critical systems might lead to failures, which are scarcely predicted in the design phase as ANNs are highly parallel and their parameters are hardly interpretable. Here we develop and evaluate a novel symbolic software verification framework based on incremental bounded model checking (BMC) to check for adversarial cases and coverage methods in multi-layer perceptron (MLP). We developed and evaluated a novel symbolic software verification framework based on incremental bounded model checking (BMC) to check for adversarial cases and coverage methods in a multi-layer perceptron(MLP). Besides, we developed an efficient SMT-based Context-Bounded Model Checker for Graphical Processing Units (ESBMC-GPU), to ensure the reliability of specific safety properties in which safety-critical systems can fail and make incorrect decisions at the cost of material damage or even put lives in danger. This paper marks the first symbolic verification framework to reason over ANNs implemented in CUDA. Our experimental results show that our approach implemented in ESBMC-GPU can successfully verify safety properties and covering methods in ANNs and correctly generate 28 adversarial cases in MLPs.

Index Terms—Artificial neural networks, Model checking, Artificial intelligence.

I. INTRODUCTION

Artificial Neural Networks (ANNs) have been used to perform various tasks in a safety-critical domain (e.g., medical diagnosis [1] and self-driving autonomous cars [2]), where failures such as misclassifications can put lives in danger. In particular, small noises and adversarial perturbations can lead ANNs to misclassify simple patterns. They can also cause severe damage in safety-critical systems, e.g., a red traffic light misclassified as green. If we formally verify these systems considering safety properties, then guarantees could be provided to ensure reliability and avoid errors in safety-critical ANNs. Given the current literature, there exists a gap between implementations of ANNs and safety guarantees. The application of formal methods, to ensure such safety properties, are not consolidated yet, e.g., the application of symbolic (software) model checking is still its infancy to solve this problem [3].

There exist various attempts to apply verification and validation techniques to ensure safety in ANNs. Recently, testing [4], training [5], and validation [6] have achieved impressive results. Sun et al. [4] show that concolic testing provides reasonable results for generating adversarial cases; the authors

describe how unsafe can ANNs be, but do not address actual implementations of ANN to provide formal safety guarantees. Zheng et al. [5] show how training techniques can bring noise robustness to the trained ANN, but no formal guarantees from a safety perspective are provided. Krogh and Vedelsby [6] describe validations techniques during the training phase, and how such a phase can be optimized by exploring the dataset and improving generalization capability. However, such works do not provide safety guarantees for the implementation of ANNs. Despite positive recent results obtained from them, symbolic software model checking can be seen as an alternative to ensure the safety of actual implementations of ANN.

In recent years, researchers proposed the verification of ANNs using formal methods [4], [7], [8]. Huang et al. [7] describe a formal verification framework based on Satisfiability Modulo Theories (SMT); they focus on verifying the safety classification of images with real cases perturbations such as scratches, changes in camera angle, or lighting conditions. Sun et al. [4] describe and evaluate concolic testing for ANNs; their test criteria measure how adversarial an input is concerning other images. An algorithm uses these test criteria based on linear programming, which generates test cases by perturbing a given system until adversarial cases are obtained. Katz et al. [8] describe a formal verification technique for checking properties of deep neural networks (DNNs). Their technique extended the simplex method to handle ReLU (Rectified Linear Unit) activation function, which is evaluated on the airborne collision avoidance system for unmanned aircraft (ACAS Xu). However, ANNs for image classification is problematic to verification approaches as the number of layers and neurons requires a considerable amount of mathematical computations.

Here we describe and evaluate a novel symbolic software verification framework based on incremental bounded model checking (BMC) to check covering methods and adversarial cases in ANNs using the efficient SMT-based Context-Bounded Model Checker for Graphical Processing Units (ESBMC-GPU) [9], [10]. Our approach, implemented on top of ESBMC-GPU, verifies the actual implementation of ANNs using CUDA and its deep learning primitives cuDNN [11] and cuBLAS [12]. Our work makes two significant contributions. First, we formally verify covering methods, as described by Sun et al. [4], to measure how adversarial two images can be to the ANN neurons using incremental BMC. Second, we formally check adversarial cases using instances of ANNs,

where only weights and bias descriptions are needed for our verification process. In both checks, we use incremental BMC to iteratively verify the implementation of ANNs for each unwind bound indefinitely (i.e., until it exhausts the time or memory limits). Intuitively, we aim to either find a counterexample with up to k loop unwinding or to fully unwind all loops in the actual implementations of ANNs so we can provide a correct result. Our incremental verification algorithm relies on a symbolic execution engine to increasingly unwind the loop after each iteration of the algorithm.

II. PRELIMINARIES

A. Artificial Neural Networks (ANNs)

Artificial Neural Networks (ANNs) are efficient models for statistical pattern recognition, which makes them a suitable paradigm for learning tasks [13]. Learning in ANNs is based on modifying the *synaptic* weights of their interconnected units, which will fit their values according to a given set of labeled examples or tasks, during the training phase. A training algorithm uses each input example to converge these weights to specific results. In most applications areas of ANNs, this may solve classification problems and some specific generalization for adversarial situations [14].

There exist various algorithms to train ANNs; *backpropagation* is a commonly used algorithm [14]. ANNs architectures vary in the type of layers, activation functions, number of layers, and neurons. Perceptron multi-layer [15] is a classic and commonly used ANN. In the safety verification of ANNs, some notations are important to logically represent operations. Hagan et al. [14] describe how Multi-layer Perceptrons (MLPs) calculate a neuron output. The neuron activation potential of an input $i = \{i_0, i_1, i_2, \dots, i_n\}$ will be denoted as u , which is obtained from

$$u = \sum_{j=0}^n w_j \times i_j + b. \quad (1)$$

$\mathcal{N}(i)$ or Y represents neuron output, that is, \mathcal{N} is the activation function of the activation potential u . Suppose that the activation function is *sigmoid*, \mathcal{N} would be described by

$$\mathcal{N}(u) = \frac{1}{1 + e^{-u}}. \quad (2)$$

ReLU [16], Sigmoid [17], and Gaussian [17] are activation functions commonly used in MLP. All our experimental results are obtained from MLPs with *sigmoid* activation function.

However, adversarial perturbation or small noises may lead to misclassification. In some cases, this may happen because the network tends to not recognize the data entered due to convergence problems, e.g., the vanishing gradient in *backpropagation* [13]. Although some techniques can be used to avoid such problems (e.g., gradient regularization [18]), safety-critical systems usually deal with a huge volume of complex data, which may be coded into big numbers and functions such as *sigmoid* squish. Thus, a large change in the input of the *sigmoid* function will cause a small change in the output; therefore, the derivative becomes small. For safety-critical applications, detecting where such problems may happen in

a complex network is a task that the formal verification community has been seeking to solve (e.g., [7]); our work aims to contribute to the solution of this research problem.

B. Incremental Bounded Model Checking (BMC)

We have used the ESBMC-GPU tool [9], [10], which is an extension of the Efficient SMT-Based Context-Bounded Model Checker (ESBMC) [19], aimed at verifying Graphics Processing Unit (GPU) programs written for the Compute Unified Device Architecture (CUDA) platform. ESBMC-GPU uses an operational model (OM), i.e., an abstract representation of the standard CUDA libraries, which conservatively approximates their semantics, to verify CUDA-based programs.

Our incremental BMC approach in ESBMC-GPU uses an iterative technique and verifies the implementation of ANNs for each unwind bound indefinitely or until it exhausts the time or memory limits. We aim to either find a counterexample with up to k loop unwinding or to fully unwind all loops. The algorithm relies on ESBMC's symbolic execution engine to increasingly unwind the loop after each algorithm iteration.

The approach is divided into two steps: one that tries to find property violations and one that checks if all the loops were fully unwound. When searching for property violation, ESBMC-GPU replaces all unwinding assertions (i.e., assertions to check if a loop was completely unrolled) by unwinding assumptions. Normally, this would lead to unsound behavior, however, the first step can only find property violations, thus reporting that an unwinding assertion failure is not a real bug. The next step is to check if all loops in the program were fully unrolled. This is done by checking whether all the unwinding assertions are unsatisfiable. No assertion is checked in the second step because they were already checked in the first step for the current k loop unwinding.

The algorithm also offers the option to change the granularity of the increment; the default value is 1. Note that changing the value of the increment can lead to slower verification time and might not present the shortest counterexample possible for the property violation. ESBMC-GPU also explicitly explores the possible interleavings of CUDA programs (up to the given context bound), while treats each interleaving itself symbolically. Additionally, ESBMC-GPU employs monotonic partial order reduction [20] and the two-thread analysis [21] to efficiently prune the state-space exploration. ESBMC-GPU verifies properties such as user-specified assertions, deadlocks, memory leaks, invalid pointer dereference, array out-of-bounds, and division by zero in CUDA programs.

III. INCREMENTAL BMC OF ANNS IN CUDA

Our goal is to detect adversarial cases that lead ANNs to wrong results. We have developed two verification strategies for ANNs described using weights and bias vectors. First, we check the coverage criteria of ANNs w.r.t. a set of images. Second, we verify safety aspects by checking adversarial cases of a given image known by an ANN.

A. SMT-based Safety Verification for ANNs

Our incremental verification method is based on two phases: (1) obtain the required models from real ANN programs

written in CUDA and (2) design safety properties that ensure the reliability of ANN implementations. Due to the benefits that parallel programming provides to train ANN algorithms, some tools such as TensorFlow [?] convert the ANN models written in Python to CUDA code, which thus exploits parallel and low-level C/C++ programming. All our verification models are collected from actual implementations of ANNs written in CUDA. There exist two APIs in CUDA to support neural networks and mathematical operations: *cuBLAS* [12] and *cuDNN* [11]. *cuBLAS* provides mathematical operations such as matrix multiplication and sums; these operations are typically used on the feed-forward process of neural networks [14]. *cuDNN* provides deep neural networks (DNN) primitives such as tensors operations, convolution functions, activation functions, and backward operations [14].

Here we have developed operational models (cf. Section II-B) for the *cuBLAS* and *cuDNN* libraries and integrated them into ESBMC-GPU.¹ Any ANN written in CUDA using these APIs can be verified by our method. Additionally, any user-specified assertion can be provided and then verified by ESBMC-GPU. Our operational models aim at performing the same steps that the execution of the original library does but ignoring irrelevant calls (e.g., screen-printing methods), where there exists no safety property to be checked. For the *cuBLAS* and *cuDNN* libraries, we ensure that their operational models return the same results as the original APIs. All developed modules were manually verified and exhaustively compared with the original ones to ensure the same behavior via conformance testing [22]. As an example, we can see a pseudo-code of the *cublasSgemv* operational model illustrated in Algorithm 1, which consists of multiplying matrices *A* and *B* and store its result in matrix *C*.

Algorithm 1 cublasSgemv

```

1: sum ← 0
2: x ← 0
3: y ← 0
4: z ← 0
5: for x < k do
6:   for y < i do
7:     for z < j do
8:       sum ← A[x][z] * B[y][x] + sum
9:       x ++
10:      C[y][z] = sum
11:      y ++
12:      z ++

```

B. Verification of Covering Methods

Covering methods [4] are based on Modified Condition/Decision Coverage (MC/DC) [23], which is a method applied to ensure adequate testing for safety-critical software; in our symbolic verification framework, MC/DC represents conditions and decisions of ANNs. Conditions are neurons in the previously layer and decisions are neurons of the following layer. Covering methods are used to measure how adversarial are two images w.r.t. the ANN neurons. There exist four covering methods available in the literature [4]: *Sign-Sign*

Cover (SS Cover), *Distance-Sign Cover* (DS Cover), *Sign-Value Cover* (SV Cover), *Distance-Value Cover* (DV Cover). Each covering method is represented as a property (assertion) within our verification framework. In particular, each property specifies that an image set must cause that all the neurons present in the ANN are covered by the covering method. Our algorithm uses each covering method to evaluate whether the adversarial behavior of a set of pair of images reaches a certain percentage of all containing neurons on the ANN.

In covering methods, neurons are indexed as n_{kl} , where k represents its sort on the layer l , inputs are defined as x_1 and x_2 . A signal-change denoted as sc occurs when the activation potential of a certain neuron has its sign function changed by two different inputs. Sign function can be described as:

$$sign(i) = \begin{cases} 1, & \text{if } i \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

Covering methods consist of some logical expressions: $sc(n_{k,l}, x_1, x_2)$ is denoted if the activation potential u of neuron $n_{k,l}$ has its signals changed, this is, $sign(u(x_1)) \neq sign(u(x_2))$. A value-change or vc occurs when the activation potential of a determined neuron represents a certain value change w.r.t. some metric h and no signal change has occurred, i.e., $h(u(x_1), u(x_2)) = true$ and $\neg sc(n_{k,l}, x_1, x_2)$, where h can be a rate function, e.g., $h(a, b) = \frac{a}{b} \geq d$, (d represents a real number that limits the change value by h); a distance-change or dc occurs when all containing neurons in a certain layer have no signal-change and its values represent some value change. We denote $dc(h, k, x_1, x_2)$ if $\neg sc(n_{k,l}, x_1, x_2)$ for all neurons in layer k and $h(k, x_1, x_2) \geq d$. As defined in value-change, function h can be instantiated as any norm-based distance function and d can be any real number that limits the distance function. A neuron pair is denoted as $\alpha = (n_{k,i}, n_{k+1,j})$ and two inputs are denoted as x_1 and x_2 . The covering methods are formally described as follows:

Sign Sign Cover (SS-Cover):

$sc(n_{k,i}, x_1, x_2);$
 $\neg sc(n_{k,l}, x_1, x_2) \forall n_{k,l} \in L_k \setminus \{i\};$
 $sc(n_{k+1,j}, x_1, x_2);$

Distance Sign Cover (DS-Cover):

$dc(h, k, x_1, x_2);$
 $sc(n_{k+1,l}, x_1, x_2);$

Sign Value Cover (SV-Cover):

$sc(n_{k,i}, x_1, x_2);$
 $\neg sc(n_{k,l}, x_1, x_2) \forall n_{k,l} \in L_k \setminus \{i\};$
 $vc(g, n_{k+1,j}, x_1, x_2);$

Distance Value Cover (DV-Cover):

$dc(h, k, x_1, x_2);$
 $vc(g, n_{k+1,j}, x_1, x_2);$

Some examples of covering methods can be seen in Table I. A multilayer perceptron that uses sigmoid function is illustrated in Fig. 1.

¹The full implementation of our OMs are available at <https://github.com/LuizHenriqueSena/ESBMC-GPU/tree/master/cpp/library>




Fig. 1: ANN instantiated.

Ex	Input	$n_{1,1}$	$n_{1,2}$	$n_{1,3}$	$n_{2,1}$	$n_{2,2}$	$n_{3,1}$
Ex1	(1, -3)	-1,3	-1,80	-0,50	-0,79	-1,370	-1,417
Ex2	(1, -1)	-0,3	-0,40	0,10	0,51	0,090	0,353
Ex3	(1, -1.2)	-0,4	-0,54	0,04	0,38	-0,056	0,176
Ex4	(1, -7)	-3,3	-4,60	-1,70s	-3,39	-4,290	-4,957

TABLE I: ANN instantiated examples and covering methods.

In Table I, Ex1 and Ex2 instances have the pair $\alpha = \{n_{3,1}, n_{1,2}\}$ covered by SS-Cover. Here only neuron $n_{3,1}$ has its signal changed in layer 1. This means that $sc(n_{3,1}, x_1, x_2)$ and $\neg sc(n_{1,l}, x_1, x_2) \forall p_{k,l} \in P_k \setminus \{3\}$ are *true*. Since only one neuron has its signal changed in layer 1, any neuron in layer 2 with its signal changed will make a pair covered by SS-Cover. In this case, if the first neuron of layer 2 ($n_{1,2}$) has its value changed, then the pair $\alpha = \{n_{3,1}, n_{1,2}\}$ is SS-Covered.

Ex2 and Ex3 are examples of DS-Cover. In Table I, there exist no change values in layer 1; it means that for some metric h (euclidian distance, gaussian), the two instances of layer 1 have a distance change and it makes $dc(h, 1, x_1, x_2)$ *true*. The distance change in layer 1 implies a signal change of neuron $n_{2,2}$; it makes the neuron pair $\alpha = \{n_{1,i}, n_{2,1}\}$ formed by any neuron in layer 1 covered by DS-Cover.

Ex2 and Ex3 contain a pair covered by SV-Cover. Similar to SS-Cover, there exists only one neuron with signal change in layer 2; it implies a value change on neuron $n_{3,1}$, which means that the neuron pair $\alpha = \{n_{2,2}, n_{3,1}\}$ is SV-Covered. Finally, Ex1 and Ex 4 are examples of DV-Cover. In this case, there exists no signal change in any layer of the ANN instances. However, if there exists a metric that makes a distance change *true* in layer 1 or 2 and if there exists another metric, which makes a value change *true* in any neuron of layer 2 or 3, then we have a DV-Covered pair.

The ANN properties to be checked by ESBMC-GPU generate the same literal to all four covering methods. The neuron covered by one of all four covering methods must be equal or greater or equal than a percentage P of all containing neurons in the ANN. In particular, these properties generate literals $l_{covered_neuron}$ with the goal of representing the validity of the covering method for the ANN, according to four constraints:

For SS-Cover:

$$l_{covered_neuron} \Leftrightarrow \left(\frac{\sum_{i,j} ss(n_{i,k}, x_1, x_2)}{N} \geq P \right). \quad (3)$$

The other three constraints: DS-Cover, SV-Cover and DV-Cover follow almost the same literal expression described in (3). The only difference will be the summation of $ds(n_{i,k}, x_1, x_2, h)$, $sv(n_{i,k}, x_1, x_2, g)$ and $ss(n_{i,k}, x_1, x_2, g, h)$ instead of $ss(n_{i,k}, x_1, x_2)$, respectively.

C. Verification of Adversarial Case

Our verification algorithm to check adversarial cases is called `checkNN`, which is implemented in CUDA; this verification algorithm checks the ANN code by focusing on the instance of the ANN. In particular, our verification algorithm requires only weights and bias descriptions. Let us assume an image input is represented by I , and m and n represent its size. The dataset will be denoted as D ; M is the universe of all possible images with size $m \times n$. Let $\delta : \mathcal{M}^{m \times n} \times \mathcal{M}^{m \times n} \rightarrow \mathbb{R}$ be an euclidian distance operator defined as follows:

$$\delta(P, Q) = \sqrt{\sum_{i=0}^n (p_i - q_i)^2}. \quad (4)$$

Suppose that any image can be casted to a vector and its values are normalized. $P = (p_1, p_2, p_3, \dots, p_n)$ and $Q = (q_1, q_2, q_3, \dots, q_n) \in \mathcal{R}^N$. Suppose that P and Q are 5×5 images, i.e., both vectors length are 25 and they are represented by the images ‘‘A’’ and ‘‘O’’, the euclidian distance equation described by Eq. (4), δ will return 2.449.

Denoting I^d as an image that belongs to the dataset, all sets can be represented as follows:

$$I^d \in \mathcal{D}^{m \times n}, I \in \mathbb{M}^{m \times n}, \mathcal{D}^{m \times n} \subseteq \mathbb{M}^{m \times n}. \quad (5)$$

Our verification algorithm `checkNN` requires three inputs: the first input is an ANN with its weights and bias descriptions; the second one is the image to be checked; and the third one is a parameter that limits the proximity of the adversarial cases. Our safety property is specified using two intrinsic operators supported by ESBMC-GPU: `assume` and `assert`. ESBMC-GPU assumes that a function call `assume(expression)` has the following meaning: if ‘‘expression’’ is evaluated to ‘‘0’’, then the function loops forever; otherwise the function returns (no side effects). In particular, the desired image to be checked can be represented by an `assume`, which is formally described as the restriction R :

$$R = \{I \in \mathbb{M}^{mn} \mid \delta(I^d, \mathcal{I}) \leq b\}. \quad (6)$$

Eq. (6) denotes that $\mathcal{I}^d \in \mathcal{M}^{m \times n}$, which will be compared with a non-deterministic image \mathcal{I} until the euclidian distance $\delta(\mathcal{I}^d, \mathcal{I})$ described by Eq. (4) is less than or equal to the parameter b . After a non-deterministic image \mathcal{I} is obtained, a safety property is checked by ESBMC-GPU, which is represented by an `assert` statement formally described as

$$\mathcal{Y}^d = \{N(\mathcal{I}), \forall \mathcal{I} \in \mathcal{R}\}. \quad (7)$$

In Eq. (7), \mathcal{Y}^d represents any ANN’s output mapped by the input \mathcal{I}^d on the given dataset. The function N represents

the ANN function $\mathcal{R}^m \rightarrow \mathcal{R}^n$. The property is taken as the negative of Eq. (7). If the output obtained by function N and the non-deterministic image is different from the mapped output, then the property is violated and a counterexample is produced. The property described by Eq. (7) generates the literal $l_{image_misclassified}$. A classification is obtained from the output values represented by the neurons of the last layer. A reference value is denoted by the variable V , a desired classification is denoted by variable D , which represents the neuron position of the output layer and i represents any other neuron position different than D . The literal $l_{image_misclassified}$ represents the validity of the original image classification, according to the constraint:

$$l_{image_misclassified} \Leftrightarrow (n_{3,D} < V) \wedge (n_{3,i} > V). \quad (8)$$

IV. EXPERIMENTAL EVALUATION

A. Description of the Benchmarks

Our evaluation employs a new character pattern recognition benchmark, which was developed by our research group with the goal of performing conformance testing [22] over our *cuBLAS* and *cuDNN* operational models (cf. Section III-A)² In particular, each experiment uses an ANN that solves the problem of vocalic pattern recognition in 5×5 images as input. The ANN was trained by the back-propagation algorithm [13]. Our dataset is composed by 100 correct vocalics with noise images and 100 non vocalic images. All vocalics are illustrated in Fig. 2.



Fig. 2: Vocalic images in benchmarks.

We conducted the experimental evaluation on a 8-core 3.40 GHz Intel Core i7 with 24 GB of RAM and Linux OS. We use CUDA v9.0, cuDNN v5.0, cuBLAS v10.1, and ESBMC-GPU v2.0.³ All presented execution times are actually CPU times, i.e., only the elapsed time periods spent in the allocated CPUs, which was measured with the times system call (POSIX system). All adversarial cases were obtained by executing the following command, which is wrapped in a shell script to iteratively unwind the program: `esbmc-gpu verifyfnn.c -I <path-to-OM> --force-malloc-success --no-div-by-zero-check --no-pointer-check --no-bounds-check --incremental-bmc`.

1) *Availability of Data and Tools*: All benchmarks, tools, and results associated with the current evaluation, are available for downloading at <https://github.com/ssvlab/ssvlab.github.io/tree/master/gpu/benchmarks>.

²Available at <https://github.com/LuizHenriqueSena/ESBMC-GPU/tree/master/cpp/library>

³Available at <http://gpu.esbmc.org/>

B. Objectives

Using the benchmarks given in Section IV-A, our evaluation has the following two experimental goals:

- EG1 (**Covering Methods**) Evaluate the performance and correctness of our symbolic verification algorithms to check all four covering methods (cf. Section III-B).
- EG2 (**Adversarial Cases**) Evaluate the performance and correctness of our verification algorithm `checkNN` (cf. Section III-C) to verify adversarial cases obtained from changing input images and parameter proximity.

C. Results

1) *Covering Methods*: in covering methods experiments, ESBMC-GPU was able to correctly verify all four methods: SS-Cover, DS-Cover, SV-Cover, and DV-Cover. The verification time did not take longer than a few seconds for checking how adversarial are two images for the ANN neurons. The fast verification was expected since all covering methods have only deterministic inputs. These results successfully answer **EG1**: ESBMC-GPU achieves reasonable performance to correctly verify all covering methods. The verified properties imply that 80% of all neurons on the ANN must be covered by the set of images. In particular, the set of images employed was the dataset used during the training phase; the tool output of all benchmarks correctly returned that covered neurons were not greater than 80%. As a result, the dataset was unable to provide 80% of neuron coverage by any covering method. Concerning this property, the average execution time of all four covering methods, when applied to a set of 200 images, is around 20 minutes; additionally, ESBMC-GPU does not generate test cases based on covering methods yet, which we leave for future work. The two generated ANN instances of the two inputs “U” and noise “U” illustrated in Fig. 3 can be seen in Table II.

Neuron	Image "U"	Noise "U"
$n_{1,1}$	-1.885322	4.619613
$n_{1,2}$	8.775419	9.796190
$n_{1,3}$	2.959348	5.743809
$n_{1,4}$	10.424796	4.046428
$n_{1,5}$	8.172012	14.466885
$n_{2,1}$	-3.863095	-9.308636
$n_{2,2}$	5.328067	5.263461
$n_{2,3}$	-3.770385	-5.705760
$n_{2,4}$	0.574238	-2.029373
$n_{3,1}$	-6.707186	-7.149290
$n_{3,2}$	-15.815082	-17.246468
$n_{3,3}$	-10.060704	-13.074245
$n_{3,4}$	-9.688183	-4.868999
$n_{3,5}$	-0.555885	3.355738

TABLE II: ANN instances of two inputs.

According to the SS-Cover equation described in Section III and the literal described by Eq. (3), we have two pairs of neuron: $\alpha_1 = \{n_{1,1}, n_{2,4}\}$ and $\alpha_2 = \{n_{2,4}, n_{3,5}\}$ SS-Covered by the two images. Only 3 of 14 neurons are SS-Covered, that is, only 21% of the neurons are covered by the covering method. The literal described by Eq. (3) specifies that the neuron coverage must be greater than a percentage P . All

benchmarks were run with the value of P set to 80%, which means that for these two inputs the property fails.



Fig. 3: Vocalics U and U with some noise.

2) *Adversarial Cases*: in adversarial cases, we executed 14 experiments, where the average verification time was approx. 19 hours. Table III shows the verification time and parameter values of our experiments. Here, *benchmark* is the experiment identifier, *image* represents the desired image, where we intend to have the classification verified, γ represents a limit of proximity for each benchmark, and *verification time* is the execution time taken by ESBMC-GPU measured in hours using times system call (POSIX system).

Benchmark	Image	γ	Verification Time (hours)
Ex1	Vocalic O	0.5	1
Ex2	Vocalic O	1.5	4
Ex3	Vocalic O	2.5	8
Ex4	Vocalic O	3.5	6
Ex5	Vocalic E	0.5	25
Ex6	Vocalic E	0.7	25
Ex7	Vocalic E	1.5	14
Ex8	Vocalic E	3.0	12
Ex9	Vocalic U	0.3	6
Ex10	Vocalic U	0.5	5
Ex11	Vocalic U	1.0	20
Ex12	Vocalic U	1.5	19
Ex13	Vocalic A	1.0	63
Ex14	Vocalic A	1.5	63

TABLE III: Verification time and proximity parameter relation.

Incremental BMC approach has not been so efficient w.r.t. verification time since we verify the actual implementation of ANNs in CUDA using the logics QF_AUFBV from the SMT standard [24], The QF_AUFBV logic represents quantifier-free formulae that are built over bit-vectors and arrays with free sort and function symbols, but with the restriction that all array terms have the following structure (array (bit-vector $i[w_1]$) (bit-vector $v[w_2]$)), where i is the index with bit-width w_1 and v is the value with bit-width w_2 .

Although our ANN benchmarks contain 3 layers, ESBMC-GPU took longer than other existing tools, which have previously reported experimental results with larger ANNs [7]. In particular, DLV (Deep Learning Verification) [7] has obtained adversarial cases of ANNs with 12 layers, ranging from a few seconds to 20 minutes. To check refinement by layer, DLV uses the theory of linear real arithmetic with existential and universal quantifiers, and for verification within a layer (0-variation), DLV uses the same theory but without universal quantification.

In prior work, Cordeiro et al. [25] reported that, although verification conditions are solved faster using the theory of linear real arithmetic (since the result of the analysis is independent from the actual binary representation), the theory of bit-vector allows the encoding of bit-level operators more accurately, which is inherently present in the implementation

of ANNs. These results partially answer **EG2**: ESBMC-GPU can correctly produce adversarial cases, which are confirmed by graphical inspection in MATLAB, but our verification time is high due to the bit-accurate precision of our verification model. Some of the adversarial cases produced by ESBMC-GPU are shown in Fig. 4.

Unfortunately, a direct performance comparison using the same ANN with DLV was not possible due to some compatibility issues. Note that DeepConcolic [4] and DLV [7] perform verification and testing only on ANNs trained by classics dataset as MNIST; both tools do not support the verification and testing on general ANNs. For future work, we will adapt our tool to support the same models as the ones handled by DeepConcolic and DLV.

D. Threats to validity

Although all adversarial cases obtained from the benchmarks are real, which we confirmed by executing our validation scripts in MATLAB, our algorithm uses a lookup table to compute the *sigmoid* activation function; even with a good resolution, a lookup table will always contain more value errors than a computed function, which could potentially produce a false adversarial case. Additionally, our new character pattern recognition benchmark is probably an easy ANN to find adversarial cases, since the ANN training was performed by a small dataset, which contains only 200 images. This kind of training probably generates ANN with a low-level of safety classification. Another identified threat is that our verification framework is not flexible enough; any change to ANN layer number might not be an easy task since our verification algorithms based on incremental BMC were developed using *cuBLAS* and *cuDNN* primitives; these primitives are not so flexible to perform such changes.

V. RELATED WORK

Our ultimate goal is to formally ensure safety for applications that are based on Artificial Intelligence (AI), as described by Amodei et al. [26]. In particular, the potential impact of intelligent systems performing tasks in society and how safety guarantees are necessary to prevent damages are the main problem of safety in ANNs.

Sun et al. [4] and Huang et al. [7] have shown how weak ANNs can be if small noises are present in their inputs. They described and evaluated testing and verification approaches based on covering methods and image proximity [4] and how adversarial cases are obtained [7]. In particular, our study resembles that of Huang et al. and Su et al. [4], [7] to obtain adversarial cases. Here, if any property is violated, then a counterexample is provided; in cases of safety properties, adversarial examples will be generated via counterexample using ESBMC-GPU. In contrast to Huang et al. [7], we do not focus on generating noise in specific regions, but in every image pixel. Our approach in image proximity is influenced by Sun et al. [4] but we use incremental BMC instead of concolic testing as our verification engine. Our symbolic verification method checks safety properties on non-deterministic images with a certain distance of a given image; both image and

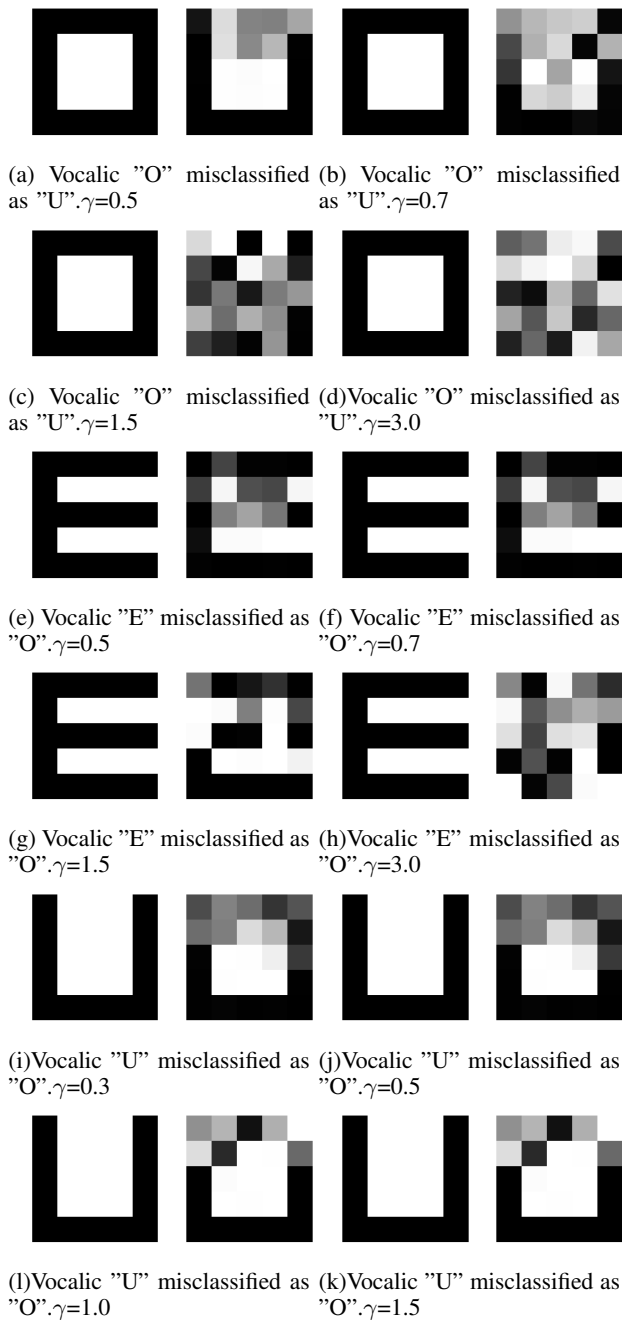


Fig. 4: Generated adversarial cases.

distance are determined by the user. Gopinath et al. [3] also describe an approach to validate ANNs using symbolic execution by translating a neural network into an imperative program. By contrast, we consider the actual implementation of ANN in CUDA and apply incremental BMC using off-the-shelf SMT solvers.

Gopinath et al. [27] presented formal techniques to extract invariants from the decision logic of ANNs. These invariants represent pre- and post-conditions, which hold when transformations of a certain type are applied to ANNs. The authors have proposed two techniques. The first one is called iterative

relaxation of decision patterns, which uses Reluplex as the decision procedures [8]. The second one is called decision-tree based invariant generation, which resembles covering methods [4]. Robustness and explainability are the core properties of this study. Applying those properties to ANNs have shown impressive experimental results. Explainability showed an important property to evaluate safety in ANNs; the core idea is to obtain explanations of why the adversarial case happened by observing the pattern activation behavior of a subset of neurons described by the given invariant.

Gopinath et al. [28] also proposed a novel approach for automatically identifying safe regions of labeled inputs. The core idea is to identify safe regions of labeled targets, i.e., providing a specific safety guarantee that a robust region is robust enough against adversarial perturbations concerning a target label. As the notion of safety robustness in ANNs is a strong term for many ANNs, the target robustness is the leading property. The technique works with clustering and verification. The clustering technique is used to split the dataset into a subset of inputs with the same labels; then, each cluster is verified by Reluplex [8] to provide the safety region of the target label. The tool proposed is called DeepSafe, which is evaluated on trained ANNs by the dataset MNIST and ACAS XU.

In addition to ESBMC-GPU, there exist other tools able to verify CUDA programs and each one of them uses its approach and targets specific property violations. However, given the current knowledge in software verification, ESBMC-GPU is the first verifier to check for adversarial cases and coverage methods in ANNs implemented in CUDA. For instance, GPUVerify [21] is based on synchronous, delayed visibility semantics, which focuses on detecting data race and barrier divergence, while reducing kernel verification procedures for the analysis of sequential programs. GPU+KLEE (GKLEE) [29], in turn, is a concrete and symbolic execution tool, which considers both *kernels* and *main* functions, while checking deadlocks, memory coalescing, data race, warp divergence, and compilation level issues. Also, Concurrency Intermediate Verification Language (CIVL) [30], a framework for static analysis and concurrent program verification, uses abstract syntax tree and partial order reduction to detect user-specified assertions, deadlocks, memory leaks, invalid pointer dereference, array out-of-bounds, and division by zero.

Our approach implemented on top of ESBMC-GPU has some similarities with other techniques described here, e.g., covering methods proposed by Sun et al. [4], model checking to solve adversarial cases proposed by Huang et al. [7]. However, the main contribution is our requirements and how we handle the actual implementations of ANNs. To run our proposed safety verification, only the ANNs with weights and bias descriptors and the desired input of the dataset is required. For tools such as DeepConcolic [4] and DLV [7], obtaining adversarial cases or safety guarantees for different ANNs is not an easy task due to the focus given to the famous datasets as MNIST [31] or CIFAR-10 [32] during the tool development. In our proposed approach, there exists no need for providing specific datasets, but only the desired dataset sample to be

verified. Besides these requirements, the user must know how cuDNN [11] deals with ANNs.

VI. CONCLUSIONS

We have described and evaluated two approaches for verifying ANNs: one to check for adversarial cases and another one to check for coverage methods in MLP. In particular, our verification method was able to find adversarial cases for different input images and proximity parameter values. Despite a high verification time in some benchmarks due to bit-accurate verification, the average time is reasonable since our approach exhaustively verifies all possible adversarial cases. In other approaches, obtaining adversarial cases of even bigger ANNs tends to be faster, but noises are not fully explored.

Our verification of covering methods was able to verify our dataset correctly and has shown to be effective during our experiments. All OMs developed here are integrated into ESBMC-GPU. They can be used to verify other programs that use the *cuBLAS* and *cuDNN* APIs. Other existing approaches (e.g., DeepConcolic and DLV) perform verification and testing only on ANNs trained by the classics dataset as MNIST, i.e., they do not offer much flexibility to perform verification and testing on general ANNs.

Future work aims to implement further techniques (e.g., invariant inference [33]) to prune the state space exploration, by taking into account design aspects of ANNs. We will also investigate fault localization and repair techniques [34] to identify and fix the root cause of incorrect ANN implementations, to make them robust against small noises in the ANN inputs.

Acknowledgement. This research, according to Article 48 of Decree n. 6.008/2006, was funded by Samsung Electronics of Amazonia Ltda, under the terms of Federal Law n. 8.387/1991, through agreement n. 004, signed with CETELI/UFAM.

REFERENCES

- [1] F. Amato, A. López, E. M. Peña-Méndez, P. Vañhara, A. Hampl, and J. Havel, "Artificial neural networks in medical diagnosis," 2013.
- [2] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [3] D. Gopinath, K. Wang, M. Zhang, C. S. Pasareanu, and S. Khurshid, "Symbolic execution for deep neural networks," *CoRR*, vol. abs/1807.10439, 2018. [Online]. Available: <http://arxiv.org/abs/1807.10439>
- [4] Y. Sun, X. Huang, and D. Kroening, "Testing deep neural networks," *arXiv preprint arXiv:1803.04792*, 2018.
- [5] S. Zheng, Y. Song, T. Leung, and I. Goodfellow, "Improving the robustness of deep neural networks via stability training," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4480–4488.
- [6] A. Krogh and J. Vedelsby, "Neural network ensembles, cross validation, and active learning," in *Advances in neural information processing systems*, 1995, pp. 231–238.
- [7] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu, "Safety verification of deep neural networks," in *International Conference on Computer Aided Verification*. Springer, 2017, pp. 3–29.
- [8] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient smt solver for verifying deep neural networks," in *International Conference on Computer Aided Verification*. Springer, 2017, pp. 97–117.
- [9] P. A. Pereira, H. F. Albuquerque, I. da Silva, H. Marques, F. R. Monteiro, R. Ferreira, and L. C. Cordeiro, "Smt-based context-bounded model checking for CUDA programs," *Concurrency and Computation: Practice and Experience*, vol. 29, no. 22, 2017.
- [10] F. R. Monteiro, E. H. d. S. Alves, I. S. Silva, H. I. Ismail, L. C. Cordeiro, and E. B. de Lima Filho, "ESBMC-GPU a context-bounded model checking tool to verify cuda programs," *Science of Computer Programming*, vol. 152, pp. 63–69, 2018.
- [11] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, "cudnn: Efficient primitives for deep learning," *arXiv preprint arXiv:1410.0759*, 2014.
- [12] C. Nvidia, "Cublas library," *NVIDIA Corporation, Santa Clara, California*, vol. 15, no. 27, p. 31, 2008.
- [13] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [14] M. T. Hagan, H. B. Demuth, M. H. Beale, and O. De Jesús, *Neural network design*. Pws Pub. Boston, 1996, vol. 20.
- [15] S. S. Haykin *et al.*, *Neural networks and learning machines/Simon Haykin*. New York: Prentice Hall., 2009.
- [16] Y. Li and Y. Yuan, "Convergence analysis of two-layer neural networks with relu activation," in *Advances in Neural Information Processing Systems*, 2017, pp. 597–607.
- [17] P. Sibi, S. A. Jones, and P. Siddharth, "Analysis of different activation functions using back propagation neural networks," *Journal of Theoretical and Applied Information Technology*, vol. 47, no. 3, pp. 1264–1268, 2013.
- [18] A. S. Ros and F. Doshi-Veleza, "Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients," in *AAAI18 - Humans and AI*. AAAI, 2018, pp. 1660–1669.
- [19] M. Y. R. Gadelha, F. R. Monteiro, J. Morse, L. C. Cordeiro, B. Fischer, and D. A. Nicole, "ESBMC 5.0: an industrial-strength C model checker," in *ASE*. ACM, 2018, pp. 888–891.
- [20] V. Kahlon, C. Wang, and A. Gupta, "Monotonic partial order reduction: An optimal symbolic partial order reduction technique," in *CAV*, ser. LNCS, vol. 5643, 2009, pp. 398–413.
- [21] A. Betts, N. Chong, A. F. Donaldson, S. Qadeer, and P. Thomson, "Gpuverify: a verifier for GPU kernels," in *OOPSLA*, 2012, pp. 113–132.
- [22] M. Krichen and S. Tripakis, "Conformance testing for real-time systems," *Formal Methods in System Design*, vol. 34, no. 3, pp. 238–304, 2009.
- [23] K. J. Hayhurst, *A practical tutorial on modified condition/decision coverage*. DIANE Publishing, 2001, vol. 210876.
- [24] C. Barrett, A. Stump, C. Tinelli, S. Boehme, D. Cok, D. Deharbe, B. Dutertre, P. Fontaine, V. Ganesh, A. Griggio, J. Grundy, P. Jackson, A. Oliveras, S. Krsti?, M. Moskal, L. De Moura, R. Sebastiani, T. D. Cok, and J. Hoenicke, "The SMT-LIB Standard: Version 2.0," *Tech. Rep.*, 2010.
- [25] L. C. Cordeiro, B. Fischer, and J. Marques-Silva, "Smt-based bounded model checking for embedded ANSI-C software," *IEEE Trans. Software Eng.*, vol. 38, no. 4, pp. 957–974, 2012.
- [26] D. Amodè, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in ai safety," *arXiv preprint arXiv:1606.06565*, 2016.
- [27] D. Gopinath, A. Taly, H. Converse, and C. S. Pasareanu, "Finding invariants in deep neural networks," *CoRR*, vol. abs/1904.13215, 2019. [Online]. Available: <http://arxiv.org/abs/1904.13215>
- [28] D. Gopinath, G. Katz, C. S. Pasareanu, and C. W. Barrett, "Deepsafe: A data-driven approach for checking adversarial robustness in neural networks," *CoRR*, vol. abs/1710.00486, 2017. [Online]. Available: <http://arxiv.org/abs/1710.00486>
- [29] G. Li, P. Li, G. Sawaya, G. Gopalakrishnan, I. Ghosh, and S. P. Rajan, "GKLEE: concolic verification and test generation for gpus," in *PPOPP*, 2012, pp. 215–224.
- [30] M. Zheng, M. S. Rogers, Z. Luo, M. B. Dwyer, and S. F. Siegel, "CIVL: formal verification of parallel programs," in *ASE*, 2015, pp. 830–835.
- [31] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *CoRR*, vol. abs/1708.07747, 2017. [Online]. Available: <http://arxiv.org/abs/1708.07747>
- [32] A. Krizhevsky, V. Nair, and G. Hinton, "The cifar-10 dataset," *online: http://www.cs.toronto.edu/kriz/cifar.html*, vol. 55, 2014.
- [33] M. Y. R. Gadelha, F. R. Monteiro, L. C. Cordeiro, and D. A. Nicole, "ESBMC v6.0: Verifying C programs using k-induction and invariant inference - (competition contribution)," in *TACAS*, ser. LNCS, vol. 11429, 2019, pp. 209–213.
- [34] E. H. da S. Alves, L. C. Cordeiro, and E. B. de Lima Filho, "A method to localize faults in concurrent C programs," *Journal of Systems and Software*, vol. 132, pp. 336–352, 2017.