PIECEWISE LINEAR NEURAL NETWORK VERIFICATION: A COMPARATIVE STUDY

Anonymous authors
Paper under double-blind review

ABSTRACT

The success of Deep Learning and its potential use in many important safety-critical applications has motivated research on formal verification of Neural Network (NN) models. Despite the reputation of learned NN models to behave as black boxes and the theoretical hardness of proving their properties, researchers have been successful in verifying some classes of models by exploiting their piecewise linear structure. Unfortunately, most of these approaches test their algorithms on their own models and do not compare with other approaches. As a result, the advantages and downsides of the different algorithms are not well understood. Motivated by the need of accelerating progress in this very important area, we investigate the trade-offs of a number of different approaches based on Mixed Integer Programming, Satisfiability Modulo Theory, as well as a novel method based on the Branch-and-Bound framework. We also propose a new data set of benchmarks, in addition to a collection of previously released testcases that can be used to compare existing methods. Our analysis not only allowed a comparison to be made between different strategies, the comparison of results from different solvers also revealed implementation bugs in published methods. We expect that the availability of our benchmark and the analysis of the different approaches will allow researchers to develop and evaluate promising approaches for making progress on this important topic.

1 INTRODUCTION

Despite their success in a wide variety of applications, Deep Neural Networks have seen limited adoption in safety-critical settings. The main explanation for this lies in their reputation for being black-boxes whose behaviour can not be predicted. Current approaches to evaluate the quality of trained models mostly rely on testing using held-out data sets. However, as Edsger W. Dijkstra said (Buxton & Randell, 1970), "testing shows the presence, not the absence of bugs". If deep learning models are to be deployed to applications such as autonomous driving cars, we need to be able to enforce and verify safety-critical behaviours.

A particularly illustrative instance of the limit of our understanding of the behaviour of learned models lies in the existence of adversarial examples (Szegedy et al., 2014): small perturbations, imperceptible to the human eye, are capable of significantly modifying the predictions generated by a network, despite it performing well on its test set. Several methods have been proposed (Gu & Rigazio, 2015; Goodfellow et al., 2015) to make networks more robust to those perturbed inputs. It is however not clear if those methods are successful at reducing the number of adversarial examples or if they are just capable of reducing the number of adversarial examples that currently known methods can generate. The only way to know this is to measure the size of the region around training samples guaranteed to not contain adversarial examples.

Some researchers have indeed tried to use formal methods to obtain guarantees like the one mentioned above. To the best of our knowledge, Zakrzewski (2001) was the first to propose a method to verify simple, one hidden layer neural networks, but only recently (Katz et al., 2017a; Narodytska et al., 2017; Cheng et al., 2017c) were researchers able to work with non-trivial models by taking advantage of the structure of ReLU-based networks. Even then, these works are not scalable to the large networks encountered in most real world problems.

A major roadblock in the area has been the lack of any analysis of the success and failure modes of the proposed approaches. To remedy this problem, we gather a data set of test cases based on existing literature and parametrically explore the space of possible architectures. We use it to evaluate different published approaches, implementing them ourselves where no public version was available, and generate the first experimental comparison of published tools. In addition, we introduce a general formalism for the problem, showing possible directions for improvement, as well as a new method showing significantly better performance on practical scenarios. Additionally, our comparison also revealed bugs in some publicly available NN verification software, made evident by contradictions in the results of the various methods.

1.1 Problem specification

We now specify the problem of formal verification of neural networks. Given a network that implements a function $y = f(x)$, a bounded input domain $\mathcal{C}$ and a property P, we want to prove that

$$\mathbf{x} \in \mathcal{C}, \quad \mathbf{y} = f(\mathbf{x}) \implies P(\mathbf{y}). \tag{1}$$

A toy-example of such a problem is given in Figure 1. On the domain $\mathcal{C} = [-2; 2] \times [-2; 2]$, we want to prove that the output y of the one hidden-layer network, always satisfy the property $P(y) \triangleq [y > -5]$. We will use this as a running example to explain the methods used for comparison in our experiments.

In this paper, we are going to focus on Piecewise-Linear Neural Networks (PL-NN), that is, networks for which we can decompose $\mathcal{C}$ into a set of polyhedra $\mathcal{C}_i$ such that $\mathcal{C} = \cup_i \mathcal{C}_i$, and the restriction of f to $\mathcal{C}_i$ is a linear function for all i. While this prevents us from including networks that use activation functions such as sigmoid or tanh, PL-NNs allow the use of linear transformations such as fully-connected or convolutional layers, pooling units such as MaxPooling and activation functions such as ReLUs and its various extensions such as Leaky ReLU or PReLU (He et al., 2015). In other words, PL-NNs represent the majority of networks used in practice. Note that as the networks to be verified are in their test-time state, layers such as Batch-Normalization still preserve piecewise linearity.

The properties that we are going to consider are Boolean formulas over linear inequalities. Although the formulas in (1) define the property to be a function of the ouptut $\mathbf{y}$, we have no loss of generality: any property involving more variables could be expressed as a function over the output of a different network, including all additional variables in its output.

The formalism introduced in (1) imposes the property to be defined only over the output. It is however not leading to any loss of generality as it would be possible to modify the network and make it "pass-through" inputs or hidden units so as to make them part of the output of the network.

As the problem of PL-NN verification has been shown to be NP-complete using a reduction from 3-SAT (Katz et al., 2017a), it is unlikely that any polynomial time algorithm will exist. Therefore, experimental comparison remains the only approach possible to evaluate the relative advantages of different methods, which we propose to do.

2 Related works

We start by briefly presenting related work that we don't include in our comparison as they aren't capable of performing general verification of problem of the form of Problem 1.

Zakrzewski (2001) and Hein & Andriushchenko (2017) propose methods based on the second derivatives of the functions expressed by the networks to guarantee that the output of the network doesn't change too much around points. However, this requires the additional assumption that all layers of the networks are twice differentiable, which PL-NN don't satisfy. At the other end of the spectrum, Narodytska et al. (2017) and Cheng et al. (2017c) use SAT solvers to propose verification methods for the specialised case of Binarized Neural Networks (Hubara et al., 2016). These methods, however, don't translate to PL-NNs.

In addition to exact formal methods, some other approaches were proposed that don't provide complete verification. Pulina & Tacchella (2010) proposed a method for verification of networks with sigmoid activation function. The necessary approximation to the smooth non-linearities lead to the incapability of their method to return a decision for certain problems. Bastani et al. (2016) studied

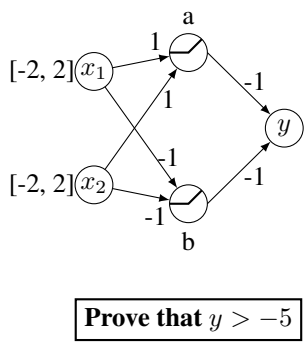

Figure 1: Example Neural Network. We attempt to prove the property that the network output is always greater than -5

Step	x_1	x_2	a_{in}	a_{out}	b_{in}	b_{out}	y
1	0	0	0	0	0	0	0
	Fix linear constraints						
	0	0	0	1	0	4	-5
2	0	0	0	1	0	4	-5
	Fix a ReLU						
	0	0	0	1	4	4	-5
3	0	0	0	1	4	4	-5
	Fix linear constraints						
	-2	-2	-4	1	4	4	-5
	...						

Table 1: Evolution of the Reluplex algorithm. Red cells corresponds to value violating Linear constraints, and orange cells corresponds to value violating ReLU constraints. Resolution of violation of linear constraints are prioritised.

PL-NN in the context of obtaining robustness guarantees against adversarial examples but added additional assumptions for scalability reasons, limiting the domain considered to the set of points sharing the same activation pattern than a reference point, effectively circumventing all of the non-linearities of the network. Xiang et al. (2017) and Huang et al. (2017) both rely on discretisation and perform layer by layer analysis to obtain guarantees over the output, but Huang et al. (2017) require the user to specify a family of possible changes and Xiang et al. (2017) over-approximates the output space of each layer, leading to the existence of undecidable properties, making both inadequate for our comparison.

3 Verification Methods

The methods we involve in our comparison all leverage the piecewise-linear structure of PL-NN to make the problem more tractable. We use the network of Figure 1 throughout as an illustrative example. All the methods we compare follow the same general principle: given a property to prove, they attempt to discover a counterexample that would make the property false. This is accomplished is by defining a set of variables corresponding to the inputs, hidden units and output of the network, as well as the set of constraints that a counterexample would satisfy. For the network of Figure 1, the variables would be $\{x_1, x_2, a_{\text{in}}, a_{\text{out}}, b_{\text{in}}, b_{\text{out}}, y\}$ and the set of constraints would be:

$$-2 \leq x_1 \leq 2 \qquad -2 \leq x_2 \leq 2 \qquad y \leq -5 \tag{2a}$$

$$a_{\text{in}} = x_1 + x_2 \qquad b_{\text{in}} = -x_1 - x_2 \qquad y = -a_{\text{out}} - b_{\text{out}} \tag{2b}$$

$$a_{\text{out}} = \max(a_{\text{in}}, 0) \qquad b_{\text{out}} = \max(b_{\text{in}}, 0) \tag{2c}$$

Here, a_{in} is the input value to hidden unit a, while a_{out} is the value after the ReLU. Any point satisfying all the above constraints would be a counterexample to the property, as it would imply that it is possible to drive the output to -5 or less. However, if this problem is unsatisfiable, no counterexample can exist, which implies that the property is True. We want to emphasize that our requirements go beyond saying that no counterexample could be found: it is necessary to prove that no counter-examples exists. The difficulty stems from the non-linear constraints of (2c). We will now explain how each method tackles this problem.

3.1 Mixed Integer Programming Encoding

A possible way to eliminate the non-linearities is to encode them with the help of binary variables, which transform the PL-NN verification problem into a Mixed Integer Program (MIP). Lomuscio & Maganti (2017) and Cheng et al. (2017b) advocate the use of big-M encoding to achieve this. For example, the non-linearities of equation (2c) are replaced by

$$\begin{gathered} a_{\text{out}} \geq 0 \qquad a_{\text{out}} \geq a_{\text{in}} \\ a_{\text{out}} \leq a_{\text{in}} + (1 - \delta_a) M_a \qquad a_{\text{out}} \leq \delta_a M_a \\ \delta_a \in \{0, 1\} \end{gathered} \tag{3}$$

where M_a is a sufficiently large value. The binary variable δ_a indicates which phase the ReLU is in: if $\delta_a = 0$, the ReLU is blocked and $a_{\text{out}} = 0$, else the ReLU is passing and $a_{\text{out}} = a_{\text{in}}$. The problem

remains difficult due to the integrality constraint on δ_a. We provide more details in Appendix A on how to handle MaxPooling units similarly.

This approach has some advantages. As the final problem to solve ends up being a MIP, imposing integrality constraints on the inputs comes at no additional cost. This can prove useful if some input features to the network are known to necessarily be integers (Cheng et al., 2017a). For other methods, imposing these integer constraints would not be possible: either the proof would be attempted on the relaxed version of the networks or it would have to be done for all possible combination of integer inputs.

As the solving of MIP is NP-hard, the performance of those methods is going to be dependent both on the quality of the solver used and of the encoding. Cheng et al. (2017b) proposes several methods to obtain a good encoding by picking appropriate values of M, in order for the quality of the linear relaxation of the MIP problem to be as high as possible. However, in the end, the problem is still resolved by a general purpose MIP solver and the question remains open whether a solver not knowing more about the specific structure of the problem can be efficient on challenging benchmarks.

3.2 RELUPLEX

Katz et al. (2017a) presents a procedure named Reluplex to verify properties of Neural Network containing linear functions and ReLU activation unit, functioning as an SMT solver using the splitting-on-demand framework (Barrett et al., 2006). The principle of Reluplex is to always maintain an assignment to all of the variables, even if some of the constraints are violated.

Starting from an initial assignment, it attempts to fix some violated constraints at each step. It prioritises fixing linear constraints ((2a) and (2b) in our illustrative example) using a simplex algorithm, even if it leads to violated ReLU constraints (2c). This can be seen in step 1 and 3 of the process shown in Table 1. If no solution to the problem containing only linear constraints exists, this shows that the counterexample search is unsatisfiable. Otherwise, all linear constraints are fixed and Reluplex attempts to fix one violated ReLU at a time, such as in step 2 of Table 1 (fixing the ReLU b), potentially leading to newly violated linear constraints. This process not being guaranteed to converge, non-linearities getting split too often are split into two cases. This generates two new sub-problems, each involving an additional linear constraint instead of the linear one. The first one solves the problem where $a_{\text{out}} = 0$, the second one where $a_{\text{out}} = a_{\text{in}}$. In the worst setting, the problem will be split completely over all possible combinations of activation patterns, at which point the sub-problems are simple LPs.

3.3 PLANET

Ehlers (2017a) also proposed an approach based on splitting the problems over the possible phase of the activation function. Unlike Reluplex, the proposed tool, named Planet, operates by attempting to find an assignment to the phase of the non-linearities. Reusing the notation introduced in Section 3.1, it assigns a value of True or False to each δ_a variable, verifying at each step the feasibility of the partial assignment. As opposed to Reluplex, this has the advantage of being easily extended to networks containing MaxPooling units.

In order to detect incoherent assignments (such as both a and b being in the greater than zero region of the ReLU in the example of Figure 1) faster, the author proposes a global linear approximation to a neural network. In addition to the existing linear constraints ((2a) and (2b)), the non linear constraints are approximated by sets of linear constraints representing the convex hull of the non-linearities. ReLUs are replaced by the set of equations:

$$a_{\text{out}} \geq 0 \qquad a_{\text{out}} \geq a_{\text{in}} \qquad a_{\text{out}} \leq \text{ub}_a \frac{a_{\text{in}} - \text{lb}_a}{\text{ub}_a - \text{lb}_a}, \tag{4}$$

where a_{out} and a_{in} are respectively the output and input of the ReLU, and ub_{in} and lb_{in} are pre-computed upper and lower bounds of the ReLUs input. This feasible domain is illustrated in Figure 2a. MaxPooling units are replaced by the set of constraints:

$$\forall i,\ \text{out} \geq \text{in}_i \qquad \text{out} \leq \sum_i (\text{in}_i - \text{lb}_{\text{in}_i}) + \max_i \text{lb}_{\text{in}_i}, \tag{5}$$

where in_i are the inputs to the MaxPooling unit and lb_{in_i} their lower bounds. A one dimensional cut of this is represented in Figure 2b. As a consequence, the whole network is approximated by an LP

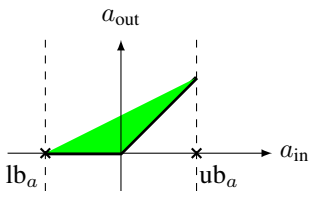

(a) Feasible domain corresponding to the set of equation of (4).

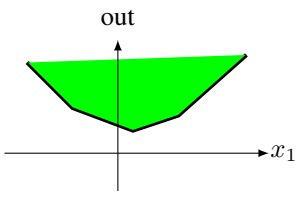

(b) Cut along one dimension of the feasible set of equations (5) where in_i are linear function of the dimension of the cut.

that can be efficiently queried to detect infeasibility or automatically deducing implied assignments to other δ_i variables. The values of the lower bounds and upper bounds necessary to define the constraints are built iteratively by optimizing the corresponding variable, based on the constraints imposed by the previous layers. Additional heuristics to make infeasibility detection and implied phase inference faster are described in the original paper (Ehlers, 2017a).

Both Reluplex and Planet rely on the splitting mechanism over ReLUs to guarantee progress but their use of it is fundamentally different: while Reluplex (Katz et al., 2017a) drives the search for satisfiability using a simplex algorithm and splits the ReLU lazily to unlock cases, Planet (Ehlers, 2017a) drives the whole search by eagerly making splits using a SAT solver and making deductions based on those. Reluplex always maintains an assignment to all variables even though it doesn't respect all constraints until the end; Planet only maintains assignments to the phase of the non-linearities. As those two approaches have never been compared on common benchmarks, it is hard to identify which is the most promising one or the specific cases in which one method outperforms the other, even though they rely on similar principles.

3.4 Branch and Bound optimization for verification

Verification as optimization We now present a different way of approaching the Neural Network verification problem. The whole satisfiability problem will be transformed into an optimization problem where the decision will be obtained by checking the sign of the minimum. We will show how any Boolean formula on linear inequalities can be encoded as additional layers at the end of the network.

Assume that the property is a simple inequality: $P(\mathbf{y}) \triangleq \mathbf{c}^T\mathbf{y} \geq b$. In that case, it is sufficient to add to the network a final fully connected layer with one output, with weight of $\mathbf{c}$ and a bias of $-b$. If the global minimum of this network is positive, it indicates that for all y reachable by the original network, we have $\mathbf{c}^T\mathbf{y} - b \geq 0 \implies \mathbf{c}^T\mathbf{y} \geq b$, and as a consequence the property is True. On the other hand, if the global minimum is negative, then it provides a counter-example to the property.

Clauses specified using OR (denoted by $\bigvee$) can be encoded by using a MaxPooling unit; if the property is $P(y) = \bigvee_i \left[\mathbf{c}_i^T\mathbf{y} \geq -b_i\right]$, this is equivalent to $\max_i \left(\mathbf{c}_i^T\mathbf{y} - b_i\right) \geq 0$. Clauses specified using AND (denoted by $\bigwedge$) can be encoded similarly: the property $P(y) = \bigwedge_i \left[\mathbf{c}_i^T\mathbf{y} \geq -b_i\right]$ is equivalent to $\min_i \left(\mathbf{c}_i^T\mathbf{y} - b_i\right) \geq 0 \iff -\left(\max_i \left(-\mathbf{c}_i^T\mathbf{y} + b_i\right)\right) \geq 0$

As a result, we can formulate any Boolean formula over linear inequalities on the output of the network as a sequence of additional layers, and the verification problem would be reduced to a global minimization problem. Aside from some specific class of NN (Amos et al., 2017), this remains a hard problem. The advantage is one of formalism, allowing us to prove the desired properties with a single procedure rather than decomposing the desired property into separate queries as was done in previous work (Katz et al., 2017a).

Finding the exact global minimum, while not necessary for verification, will have the advantage of generating a value. If this value is positive, it will correspond to the margin by which the property is satisfied. When estimating robustness to adversarial examples, existing methods choose to perform a binary search over the maximum radius guaranteeing the absence of adversarial examples. The optimization process would be a more appropriate formalism here.

Branch and Bound for Optimization Optimization algorithms such as Stochastic Gradient Descent which are the usual workhorses of Deep Learning are not appropriate for this minimization problem. Despite being capable of converging to good local minima, they have no way of guaran-

teeing whether or not a minima is global. We now present an approach to tackle this problem, based on the Branch-and-Bound method.

Algorithm 1 describes the generic form of the Branch-and-Bound method. The input domain will be repeatedly split into sub-domains (line 7), over which we will compute lower and upper bounds of the minimum of the output (lines 9-10). The best upper-bound found so far will serve as a candidate for the global minimum. Any domain whose lower bound is greater than the current global upper bound can be pruned away as it necessarily won't contain the global minimum (line 13, lines 15-17). By iteratively splitting the domains, we will be able to compute tighter lower bounds. We keep track of the global lower bound on the minimum by taking the minimum over the lower bounds of all sub-domains (line 19). When the global upper bound and the global lower bound differ by less than a small scalar ϵ (line 5), we consider that we have converged.

Algorithm 1 Branch and Bound

1: **function** BAB(net, domain, ϵ) ▷ Global minimum of net on domain to an ϵ accuracy
2: global_ub $\leftarrow \inf$
3: global_lb $\leftarrow -\inf$
4: domains $\leftarrow [(\texttt{global_lb}, \texttt{domain})]$
5: **while** global_ub − global_lb $> \epsilon$ **do**
6: $(_, \texttt{dom}) \leftarrow \texttt{pick_out}(\texttt{domains})$
7: $[\texttt{dom_1}, \ldots, \texttt{dom_s}] \leftarrow \texttt{split}(\texttt{dom})$
8: **for** $i = 1 \ldots s$ **do**
9: dom_ub $\leftarrow$ compute_upper_bound(net, dom_i)
10: dom_lb $\leftarrow$ compute_lower_bound(net, dom_i)
11: **if** dom_ub $<$ global_ub **then** ▷ If found a better upper bound
12: global_ub $\leftarrow$ dom_ub
13: domains $\leftarrow$ prune_domains(domains, global_ub) ▷ Remove domains with dom_lb $\geq$ global_ub
14: **end if**
15: **if** dom_lb $<$ global_ub **then**
16: domains $\leftarrow$ domains $+ [(\texttt{dom_lb}, \texttt{dom_i})]$
17: **end if**
18: **end for**
19: global_lb $\leftarrow \min\{\texttt{lb} \mid (\texttt{lb}, \texttt{dom}) \in \texttt{domains}\}$
20: **end while**
21: **return** global_ub
22: **end function**

The description of the verification problem as optimization and the pseudo-code of Algorithm 1 are generic and would apply to verification problems beyond the specific case of PL-NN. To obtain a practical algorithm, it is necessary to specify the elementary functions:

- pick_out(domains): Select one of the domains to branch on. Several heuristics are possible, based on the bounds previously computed on the domains or based on the size of the domains.
- split(domain): Takes as argument a domain and returns a partition of domains $[\texttt{dom_1}, \ldots, \texttt{dom_s}]$ such that $\bigcup_i \texttt{dom_i} = \texttt{domain}$ and that $(\texttt{dom_i} \cap \texttt{dom_j}) = \emptyset$, $\forall i \neq j$. Choosing the split function will define the "shape" of the domains that we are operating on, potentially making the computation of the bounds harder or easier.
- compute_{lower, upper}_bound(net, domain): Compute a (lower / upper) bound of the minimum of net over the feasible domain domain. We want the lower bound to be as high as possible (so that this whole domain can be pruned easily) and the upper bound to be as low as possible (so that we can use it to prune out other regions of the search space). If we have several approaches to compute bounds, we can employ all of them at once and only keep the tightest. In our experiments, we use the result of minimising the variable corresponding to the output of the network, subject to the constraints of the linear approximation introduced by Ehlers (2017a) as lower bound and simple random sampling

Data set	**Numbers of Properties**	**Model Architecture**
ACAS	188	5 inputs 6 layers of 50 hidden units 5 outputs
CollisionDetection	500	6 inputs 40 hidden unit layer, Maxpool 19 hidden unit layer 2 outputs
TwinStream	81	{5, 10, 25} inputs {2, 4, 5} layers of 2*{5, 10, 25} hidden units, 1 output, with a bias of {0.01, 1, 10}

Table 2: Dimensions of all the data sets

as upper bound. As the procedure is generic, any improved algorithm leading to faster or tighter bounds would directly translate to improvements to the verification process.

In practice, it is not necessary to run Algorithm 1 to convergence for verification. If a negative global upper bound is found, the corresponding input is a valid counter-example. Similarly, as soon as the global lower bound goes above zero, we know that the property is verified.

Now equipped with an understanding of the workings of each method, we are now ready to compare them.

4 Solver and Properties

As each method implements a different strategy and the worst-case analysis always indicate exponential runtimes, these approaches can only be compared experimentally. We attempt to perform verification on three data sets of verification properties and report the comparison results. The dimensions of all of the problems are given in Table 2.

For each of the data sets, we compare four methods using the same protocol. Each run is done with a timeout of two hours, and a maximum allowed memory usage of 20GB, on a single core of a machine with an i7-5930K CPU and 32GB of RAM. We will release all code and data necessary to replicate our analysis.

Success rate corresponds to the proportion of properties for which the solver returns the correct answer before timing out or being terminated for using too much memory. We compare the performances of each method and present the results separately for the cases where the properties are True or False. **SAT** means that the satisfiability problem for a counter-example was satisfiable, implying that the property was False. In this case, the runtime corresponds to the time it took before exhibiting a counter-example to the property. On the other hand, **UNSAT** corresponds to the time it took to prove that the problem was infeasible and that no counterexamples to the property could exist. As methods returning **SAT** exhibit a counterexample, disagreement between solvers can be easily resolved by evaluating the network over the counterexample and checking the property on the output. We use this criterion to establish a ground truth for each property, except for Table 5 where we know that the result is **UNSAT** by construction. We reported bugs to the original authors of each method every time such a disagreement was detected.

In the case of a timeout, the runtime for the method is counted as the maximum allowed time (7200 s), even though the actual runtime would be worse. As a result, the average runtime for methods with low success rate would be worse in practice than reported here. After fixing the Planet solver, no Out of Memory error was encountered. To give more insights into the relative performance of each solver, we count the *Number of wins*, which corresponds to how many times a solver was the fastest to solve a property. When the relative difference between the runtime of two solvers is less than 1%, we don't count any win.

Reluplex, based on the version released by the authors (Katz et al., 2017b). The tool is implemented in C++ and relies on a modified version of the GLPK library to deal with the Simplex algorithm. Note that as the tool doesn't support MaxPooling units out of the box, we automatically convert the

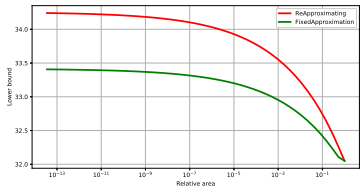

(a) Approximation on a shallow net from the CollisionDetection data set

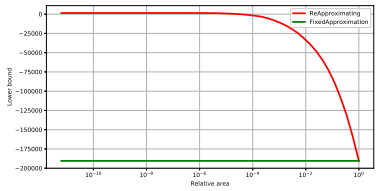

(b) Approximation on a deep net from the ACAS data set

Figure 3: Quality of the linear approximation, depending on the size of the domain. We compute the value of the lower bound on a given domain, centred around the global minimum and repeatedly shrink the size of the domain. Rebuilding completely the linear approximation at each step allows to create tighter lower-bounds thanks to better lb_a and ub_a, as opposed to using the same constraints and only changing the bounds on input variables. This effect is even more significant on deeper networks.

MaxPooling layers into a series of linear layers with ReLU activations. To do so, we decompose the element-wise maximum into a series of pairwise maximum

$$\max(x_1, x_2, x_3, x_4) = \max(\max(x_1, x_2), \max(x_3, x_4)), \tag{6}$$

and decompose the pairwise maximums as sum of ReLUs:

$$\max(x_1, x_2) = \max(x_1 - x_2, 0) + \max(x_2 - \text{lb}_{x_2}, 0) + \text{lb}_{x_2}, \tag{7}$$

where lb_{x_2} is a pre-computed lower bound of the value that x_2 can take.

Planet, based on the version released by the author (Ehlers, 2017b). The tool is implemented in C++, using GLPK to solve linear programs and a modified version of MiniSat to drive the search. We wrote some software to convert in both directions between the input format of both Reluplex and Planet. We discovered some memory issues on the original implementation, which we reported to the author and used the fixed version for all experiments.

MIP, that consists on encoding the satisfiability problem as a Mixed Integer Program, using the "big M" encoding of the non-linearity. The exact encoding of MaxPooling and ReLU can be found in Appendix A. This method is similar to the one implemented in (Cheng et al., 2017b; Lomuscio & Maganti, 2017) but due to the lack of availability of open-sourced methods, we reimplemented the approach in Python, using the Gurobi MIP solver. To choose the value of M for each ReLU, we made use of the linear approximation of Planet. This leads to better values of M than the dataflow analysis discussed by Cheng et al. (2017b). Even when compared to their other proposed heuristic for the selection of M, it has several advantages: it encompasses information from all layers of the network rather than from only a few layers; it solves simple LPs rather than MIPs.

BaB, based on the method described in Section 3.4. Our `pick_out` strategy chooses the domain that currently has the smallest lower bound. We split the domain in half across its longest edge to generate two new smaller domains as our `split` method. We generate upper bounds on the minimum by randomly sampling points on the considered domain, and minimise the linear approximation of the network proposed by (Ehlers, 2017a) as a lower bound. Our implementation is in Python and uses Gurobi to solve LPs. Note that as opposed to the approach taken by Ehlers (2017a) of building a single approximation of the network, we rebuild a new approximation for each sub-domain. This is motivated by the observation shown in Figure 3 which demonstrate the significant improvements it brings, especially for deep networks.

Having now presented the software tools we use, we know report their performance on each of our data set.

5 Analysis

5.1 Airborne Collision Avoidance System

The Airborne Collision Avoidance System (**ACAS**) data set, as released by Katz et al. (2017a) is a neural network based advisory system recommending horizontal manoeuvres for aircraft in order to

Data Set	Success Rate	Average runtime		Number of Wins	
		SAT	UNSAT	SAT	UNSAT
Reluplex	79.79%	2077.7 s	1868.3 s	3	6
Planet	46.28%	7200 s	2532.8 s	0	1
MIP	49.47%	6014.0 s	2708.2 s	3	2
BaB	83.51%	128.14 s	733.55 s	29	118

Table 3: Results on the **ACAS** data set. Note that 26 of the properties were not solved by any of the methods and are therefore not present in the SAT/UNSAT breakdowns. Reluplex and BaB boast comparable success rate but BaB finishes significantly faster, especially to find counterexamples.

Data Set	Success Rate	Average runtime		Number of Wins	
		SAT	UNSAT	SAT	UNSAT
Reluplex	99.4 %	1.14 s	1.17 s	70	30
Planet	100 %	0.50 s	0.18 s	79	275
MIP	100 %	0.66 s	0.65 s	17	11
BaB	100 %	5.27 s	29.05 s	0	0

Table 4: Results on the **CollisionDetection** data set. All solvers finished on all test cases but Reluplex erroneously classified three properties as True. For 18 properties, the difference between the two fastest methods was inferior to 1% of their runtime so we didn't count any Win.

avoid collisions, based on sensor measurements. Each of the five possible manoeuvres is assigned a score by the neural network and the action with the minimum score is chosen. In the case of complex property to prove, such as "NoAction (CoC) is the minimal score", their counter-example search is implemented as a series of four satisfiability problems: "Does there exist a point where the score for {WeakLeft, WeakRight, StrongLeft, StrongRight} is less than the score for NoAction?" while our approach discussed in Section 3.4 would be able to merge all of these satisfiability problems into a single one. For our experiments, we use the original strategy of the authors.

Results for the ACAS data are shown in Table 3. Our proposed method using Branch and Bound performs the best along all criteria. Compared to the second best method, Reluplex, it is more than an order of magnitude faster at exhibiting counter-examples and more than twice as fast at proving the correctness of True properties. On this data set which is the one with the largest number of hidden units, Planet doesn't manage to exhibit any counterexample to a property. We hypothesise that with such a high number of non-linearities, the strategy of eager assignments to phase of Planet is disadvantaged as opposed to the more lazy approach of Reluplex and BaB. All methods but BaB perform better on **UNSAT** problem than on **SAT** ones. We postulate that this is due to the relatively small dimensionality of the input, which makes random testing capable of easily getting good coverage and discovering counterexamples quickly.

5.2 COLLISIONDETECTION

In the **CollisionDetection** data set, as released by the authors of Planet, the network attempts to predict whether two vehicles with parameterized trajectories are going to collide. A total of 500 properties are extracted from problems arising from a binary search to identify the size of the region around training examples in which the prediction of the network doesn't change.

Planet is the best performing method on the set of benchmarks that accompanied the release of the tool. Table 4 shows it being the fastest method for most of the properties. Conversely, BaB becomes the worst performing method, especially on True properties. It is however important to note that all solvers finished significantly below the timeout limit, indicating that the data set isn't extremely challenging. On this data set, Reluplex classified three False properties as unsatisfiable. We evaluated the counterexamples returned by other solvers using Reluplex's code and confirmed that they were valid counterexamples. We reported the issue to the original authors of Reluplex who identified numerical instabilities as the reason for the discrepancy.

5.3 TWINSTREAM

To get a better understanding of what factors influence the performance of various methods, we propose a novel **TwinStream** data set, which we generated to have control over different hyper-

Data Set	Success Rate	Average Runtime	Number of Wins
Reluplex*	53.09%	2061.8s	1
Planet	62.96%	2878.5s	13
MIP	65.43%	2647.0s	41
BaB	28.40%	5248.1s	4

Table 5: Results on the **TwinStream** data set. Reluplex returns erroneous results for several of the properties due to numerical instabilities. Planet and MIP performs the best and Branch-and-Bound struggle on a lot of the instances.

parameters. The networks contain two separate streams, where each of the streams has the same architecture, weights, and inputs. The final layer of the network computes the difference between the outputs of the two streams, and add a positive bias term, which we will refer to as the margin. As a result, the output is always equal to the value of the final bias. On each of those networks, we attempt to prove the true property that the output of the network is positive. We generate networks with random weights using Glorot initialisation (Glorot & Bengio, 2010), and vary the depth, number of hidden units in each of the stream, number of inputs, and the value of the margin. Note that as opposed to the other two data sets, the weights aren't the result of an optimization process and therefore may not be representative of real use-cases.

For 16 out of the 27 networks having a margin of 0.01 and one of the networks having a margin of 1, Reluplex returns a SAT result although the properties are True by construction. Evaluating the returned counterexamples using Reluplex itself confirms that they are not valid counterexamples, but caused by numerical issues. We also reported these issues to the authors of Reluplex.

Surprisingly, the generic MIP solver and the Planet solver perform best on this data set. A possible explanation is that due to the heavy structure present in the problem (each non-linearity having a twin exactly equal to it in the other stream), methods explicitly representing the decision over the phase of the non-linearity inherently have the advantage.

We generate a plot showing the evolution of runtimes for each of the architectural parameters of TwinStream networks. Each curve represented corresponds to a fixed configuration of all other parameters. This allows us to assess the impact of each factor on the performance of solvers.

While MIP (Figure 4b) and Planet (Figure 4c) are relatively insensitive to changes in the number of inputs in the network, Branch-and-Bound (Figure 4a) incurs significantly higher run-times as the number of inputs to the network rises. This can be explained by the fact that MIP and Planet branch over the hidden units whose number is kept constant here, while BaB branches over the dimensions of the input. On the other hand, when studying the impact of the network depth, MIP and Planet (Figure 5b and 5c respectively) are much more sensible to the higher number of layers, unlike BaB (Figure 5a) who surprisingly performs better on deeper networks.

The analysis of Figure 6 doesn't reveal any surprising insights: higher number of non-linearities in each layer corresponds to more complex problem, which directly translates to longer runtime for all solvers. Figure 7 reveals that the expected runtime of Branch-and-Bound is more sensible to the amount by which a property is True, compared to solvers that reason directly other the phases of non-linearities. This is natural as in the case of BaB, properties with a higher margin corresponds simply to cases where it is possible to stop the counter-example search procedure faster.

6 Conclusion

The improvement of formal verification of Neural Networks represents an important challenge to be tackled before learned models can be used in safety critical applications. The lack of a shared set of benchmarks between researchers makes it hard to evaluate progress or estimate promising research directions. We gathered test cases from existing literature, proposed new ones, and evaluated the performance of published methods, which allowed us to surface issues in published tools and offer an informed view of the status of the field. In addition, we proposed a conceptually simple method that offers competitive performance on the most realistic data set available. Our method is sufficiently general to be easily improvable, as any better lower bound will directly translate to faster verification.

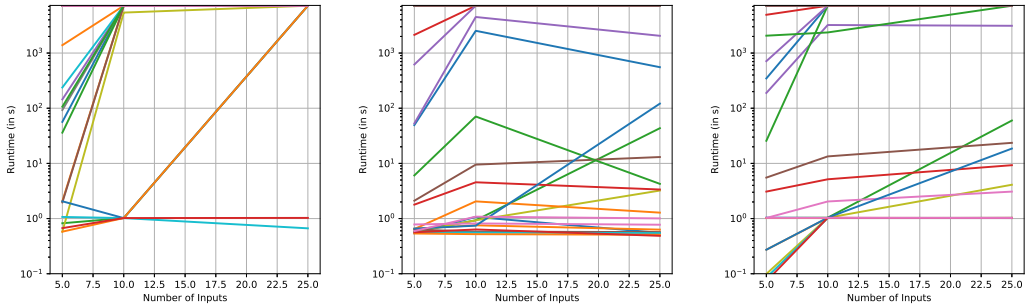

(a) Branch-and-Bound (b) MIP (c) Planet

Figure 4: Impact of the number of inputs to the network on the performance of the solvers. Branch-and-Bound performs significantly worse for input domain with a larger number of dimension, while MIP and Planet degrades more gracefully.

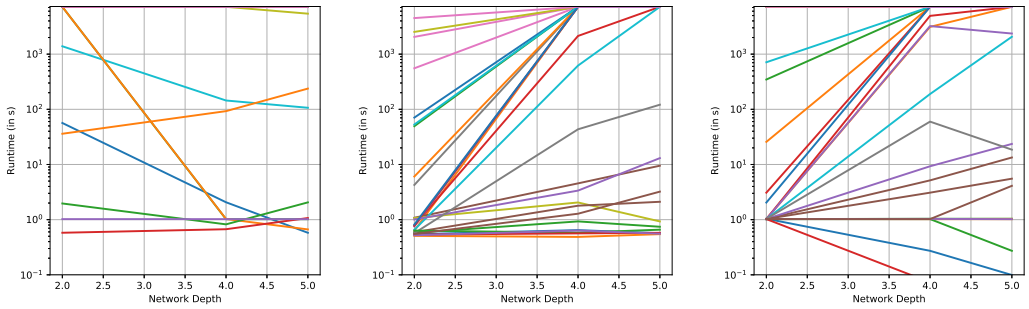

(a) Branch-and-Bound (b) MIP (c) Planet

Figure 5: Impact of the Network depth on the performance of the solvers. The deeper the networks are, the slower it is to prove a property. Surprisingly, Branch-and-Bound seems to perform better for deeper networks.

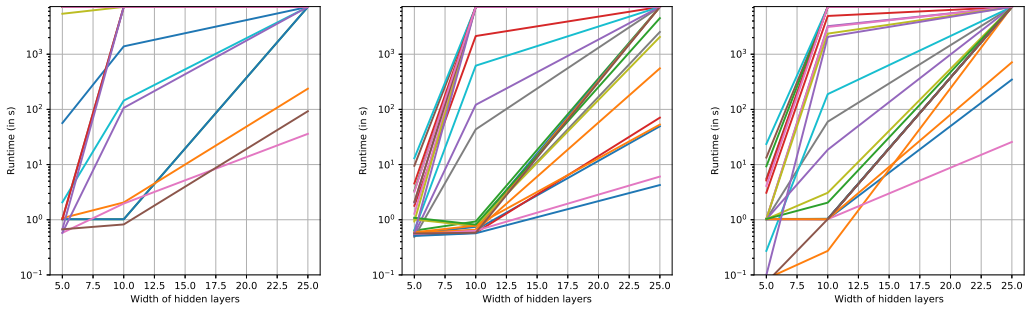

(a) Branch-and-Bound (b) MIP (c) Planet

Figure 6: Impact of the number of hidden units per layer on the performance of the solvers. More hidden units correspond to a higher number of non-linearities and consequently a harder problem to solve. All methods get worse in a similar fashion when more hidden units are introduced.

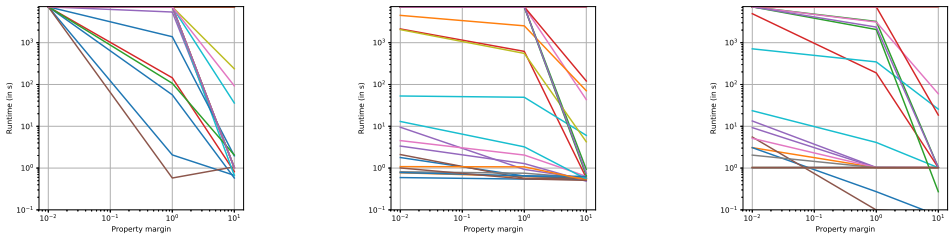

(a) Branch-and-Bound (b) MIP (c) Planet

Figure 7: Impact of the margin by which the property is True on the performance of the solvers. Branch-and-Bound is the most subject to improvement when the difficulty of the problem is relaxed, while methods based on explictly assigning values to phase of the non-linearities are less affected.

REFERENCES

Brandon Amos, Lei Xu, and J. Zico Kolter. Input convex neural networks. *ICML*, 2017.

Clark Barrett, Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Splitting on demand in sat modulo theories. *International Conference on Logic for Programming Artificial Intelligence and Reasoning*, 2006.

Osbert Bastani, Yani Ioannou, Leonidas Lampropoulos, Dimitrios Vytiniotis, Aditya Nori, and Antonio Criminisi. Measuring neural net robustness with constraints. *NIPS*, 2016.

John N Buxton and Brian Randell. *Software Engineering Techniques: Report on a Conference Sponsored by the NATO Science Committee*. NATO Science Committee, 1970.

Chih-Hong Cheng, Frederik Diehl, Yassine Hamza, Gereon Hinz, Georg Nührenberg, Markus Rickert, Harald Ruess, and Michael Troung-Le. Neural networks for safety-critical applications-challenges, experiments and perspectives. *arXiv:1709.00911*, 2017a.

Chih-Hong Cheng, Georg Nührenberg, and Harald Ruess. Maximum resilience of artificial neural networks. *Automated Technology for Verification and Analysis*, 2017b.

Chih-Hong Cheng, Georg Nührenberg, and Harald Ruess. Verification of binarized neural networks. *arXiv:1710.03107*, 2017c.

Ruediger Ehlers. Formal verification of piece-wise linear feed-forward neural networks. *Automated Technology for Verification and Analysis*, 2017a.

Ruediger Ehlers. Planet. `https://github.com/progirep/planet`, 2017b.

Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. *AISTATS*, 2010.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *ICLR*, 2015.

Shixiang Gu and Luca Rigazio. Towards deep neural network architectures robust to adversarial examples. *ICLR workshop*, 2015.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. 2015.

Matthias Hein and Maksym Andriushchenko. Formal guarantees on the robustness of a classifier against adversarial manipulation. *NIPS*, 2017.

Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. Safety verification of deep neural networks. *International Conference on Computer Aided Verification*, 2017.

Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. *NIPS*, 2016.

Guy Katz, Clark Barrett, David Dill, Kyle Julian, and Mykel Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. *CAV*, 2017a.

Guy Katz, Clark Barrett, David Dill, Kyle Julian, and Mykel Kochenderfer. Reluplex. `https://github.com/guykatzz/ReluplexCav2017`, 2017b.

Alessio Lomuscio and Lalit Maganti. An approach to reachability analysis for feed-forward relu neural networks. *arXiv:1706.07351*, 2017.

Nina Narodytska, Shiva Prasad Kasiviswanathan, Leonid Ryzhyk, Mooly Sagiv, and Toby Walsh. Verifying properties of binarized deep neural networks. *arXiv:1709.06662*, 2017.

Luca Pulina and Armando Tacchella. An abstraction-refinement approach to verification of artificial neural networks. *CAV*, 2010.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *ICLR*, 2014.

Weiming Xiang, Hoang-Dung Tran, and Taylor T Johnson. Output reachable set estimation and verification for multi-layer neural networks. *arXiv:1708.03322*, 2017.

Radosiaw R Zakrzewski. Verification of a trained neural network accuracy. *IJCNN*, 2001.

A MIP ENCODING DETAILS

We now present the encoding used to convert the non linearity of a neural network into inequality constraints involving binary variables.

In the case of a ReLU activation

$$a_{\text{out}} = \max\left(a_{\text{in}}, 0\right) \tag{8}$$

we can replace this equation containing the $\max$ non-linearity by

$$\begin{array}{cc} a_{\text{out}} \geq 0 & a_{\text{out}} \geq a_{\text{in}} \\ a_{\text{out}} \leq a_{\text{in}} + (1-\delta_a)M_a & a_{\text{out}} \leq \delta_a M_a \\ \multicolumn{2}{c}{\delta_a \in \{0,1\}} \end{array} \tag{9}$$

where M_a is a value such that M_a is an upper-bound of a_{in} and $-M_a$ is a lower-bound of a_{in}.

For MaxPooling units, we can either make the choice to decompose the MaxPooling into a series of ReLU and use their linear encoding. Another solution is to encode the MaxPooling directly. The constraint

$$\text{out} = \max_i \left(\text{in}_i\right) \tag{10}$$

can be replaced by

$$\begin{aligned} &\text{out} \geq \text{in}_i \quad \forall i \\ &\text{out} \leq \text{in}_i + (M - \text{lb}_i)(1-\delta_i) \quad \forall i \\ &\sum_i \delta_i = 1 \\ &\delta_i \in \{0,1\} \quad \forall i \end{aligned} \tag{11}$$

where M is an upper-bound on all in_i and lb_i are lower-bounds on each in_i.