

A review of abstraction methods towards verifying neural networks

FATEH BOUDARDARA and ABDERRAOUF BOUSSIF, Technological Research Institute Railenium, France
PIERRE-JEAN MEYER and MOHAMED GHAZEL, Univ Gustave Eiffel, COSYS-ESTAS, France

Neural networks as a machine learning technique are increasingly deployed in various domains. Despite their performances and their continuous improvement, the deployment of neural networks in safety-critical systems, in particular for autonomous mobility, remains restricted. This is mainly due to the lack of (formal) specifications and verification methods and tools that allow for getting sufficient confidence in the behavior of the neural network-based functions. Recent years have seen neural network verification getting more attention; and many verification methods were proposed, yet the practical applicability of these methods to real-world neural network models remains limited. The main challenge of neural network verification methods is related to the computational complexity and the large size of neural networks pertaining to complex functions. As a consequence, applying abstraction methods for neural network verification purposes is seen as a promising mean to cope with such issues. The aim of abstraction is to build an *abstract* model by *omitting* some irrelevant details or some details that are not highly impacting w.r.t some considered features. Thus, the verification process is made faster and easier while preserving, to some extent, the relevant behavior regarding the properties to be examined on the original model. In this paper, we review both the abstraction techniques for activation functions and model size reduction approaches, with a particular focus on the latter. The review primarily discusses the application of abstraction techniques on feed-forward neural networks, and explores the potential for applying abstraction to other types of neural networks. Throughout the paper, we present the main idea of each approach, and then discuss their respective advantages and limitations in details. Finally, we provide some insights and guidelines to improve the discussed methods.

CCS Concepts: • **Theory of computation** → *Abstraction; Logic and verification*; • **Computing methodologies** → **Neural networks**.

Additional Key Words and Phrases: Formal verification, neural network verification; Abstraction; Abstract interpretation;

ACM Reference Format:

Fateh Boudardara, Abderraouf Boussif, Pierre-Jean Meyer, and Mohamed Ghazel. 2022. A review of abstraction methods towards verifying neural networks. 1, 1 (September 2022), 18 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Neural Network (NN) is one of the most popular machine learning techniques [20, 36]. The use of such an approach has shown fast progress during the last decade, giving rise to a noticeable enhancement of the technique, as witnessed by its successful achievements in various domains [41]. Nowadays, applications of NNs can be encountered in a wide range of domains, such as in financial transactions, trading, forecasting and fraud detection [41, 47]. In recent years, with the advances in terms of computational performances, NNs have been widely adopted in image recognition and object detection systems [47]. Namely, they are increasingly investigated to be deployed for safety-critical applications, in particular for the design of environment monitoring and decision-making functions in autonomous vehicles and

Authors' addresses: Fateh Boudardara, fateh.boudardara@railenium.eu; Abderraouf Boussif, abderraouf.boussif@railenium.eu, Technological Research Institute Railenium, 180 rue Joseph-Louis Lagrange, Valenciennes, France, F-59308; Pierre-Jean Meyer, pierre-jean.meyer@univ-eiffel.fr; Mohamed Ghazel, mohamed.ghazel@univ-eiffel.fr, Univ Gustave Eiffel, COSYS-ESTAS, 20 rue Élisée Reclus, F-59666, Villeneuve d'Ascq, France.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM Transactions on Embedded Computing Systems

Manuscript submitted to ACM Transactions on Embedded Computing Systems

trains [62]. A software module of a safety-critical system needs to be certified before its deployment. Thus, it is required to develop methods to verify safety specifications and certify such NN-based software [25, 26].

The earliest works that deal with the verification of NN models are based encoding the model at hand as a system of linear equations, which can then be solved using off-the-shelf verification tools, namely SAT/SMT solvers [15, 30, 50] and MILP solvers [8, 12, 42, 58]. Although these methods are theoretically sound¹ and complete², they are limited to small-size neural networks due to the non-linearity of NN models. Indeed, the number of linear constraints grows exponentially with the number of neurons for which the activation functions need to be linearized, which may give rise to a state-space explosion problem. Therefore, verification methods based on over-approximation have been proposed to help mitigate this problem while preserving the soundness but not the completeness [14, 22, 32, 38, 51, 64–66, 69] (see [63] for more details). Among these techniques, abstraction methods try to ease the verification problem by abstracting the activation function using linear bounds [15] or abstract domains [18, 54, 55], or by reducing the size of the network to improve the scalability of NN verification engines. In the latter case, an abstract (or reduced) model, which is smaller and easier to verify, is generated from the original network [16, 49]; thus, instead of applying the verification method directly on the original model, the verification process can be enhanced by applying it on the reduced model.

Regarding the substantial interest in NN verification and the amount of existing methods for certifying NNs, many surveys and reviews on NN verification methods have been proposed in the literature. For instance, Leofante et al. [37] established three types of NN verification properties: equivalence, invertibility and invariance. They also provided a review of NN verification techniques based on constraints solving. Liu et al. [40] classified the existing verification methods into three basic categories: optimization, reachability and search-based verification techniques. Huang et al. [23] conducted a review about deep NN safety and trustworthiness. For NN verification, the authors distinguished between global and local properties. Regarding the guarantees of the verification technique, the survey classifies NN verification techniques into deterministic, approximative and statistical. According to [61], verification methods can be classified as geometric-based methods, MILP, SAT/SMT and optimization-based methods, even though MILP and SAT/SMT based verification methods can also be considered as particular cases of optimization techniques. Recently, Urban et al. [63] discussed the verification methods applied to machine learning. For NN verification, the authors proposed a classification of the existing methods into complete or incomplete methods with respect to the output of the verification process. Moreover, the review [63] summarizes formal verification approaches for different machine learning techniques such as support vector machine and decision trees. Another research area that is surveyed in [5] involves exploiting the sequential behavior of Recurrent Neural Networks (RNNs) to convert RNN models into automata and verify certain properties on the resulting automaton model. Although such techniques can also be seen as a form of NN abstraction into automata, their focus on recurrent neural networks places them out of the scope of the present survey where we primarily consider feed-forward neural networks.

Among all the surveys and reviews discussed above, and to the best of our knowledge, no existing work offers an overview on the abstraction methods for feed-forward NNs verification purposes. The aim of this work is to present a review on the existing activation function abstraction and model reduction methods in the literature for NN verification, and derive a critical discussion regarding these techniques. Concretely, for each presented approach we will sketch out the main idea, analyze its advantages and its drawbacks, and discuss the corresponding formal guarantees. For model reduction techniques, we will particularly highlight how each method can affect the verification process, and we will

¹Whenever the method returns that the property holds, it indeed holds on the system.

²The verification method never returns "Unknown".

105 discuss further research directions in terms of these techniques. Although this work focuses on feed-forward NNs, we
 106 also provide some perspectives on how these abstraction methods can be adjusted to support other types of NNs. It is
 107 worth noticing that in this paper, we only consider NN abstraction methods that are used for verification purposes, *i.e.*,
 108 we do not include neural networks' compression techniques such as quantization and edges pruning [21], since their
 109 goal is to build a compressed model to speed up the run-time execution [7], while preserving the model's accuracy but
 110 not necessarily its behavior, neither providing formal guarantees on the compressed model.
 111

112 The remainder of the paper is structured as follows: In Section 2, preliminary concepts and notations pertaining to
 113 neural networks are introduced, the verification problem of NNs is stated and an overview of the existing NN verification
 114 methods is provided. Section 3 reviews existing NN abstraction approaches, with a deeper focus on model reduction
 115 methods. Besides discussing the main features of the evoked techniques, some pointers to possible enhancements of the
 116 discussed methods will be provided. Finally, in Section 4 we recall the main findings through our review and outline
 117 some challenges and perspectives regarding NN abstraction.
 118
 119

120 2 BACKGROUND

121 2.1 Neural networks

122 A feed-forward neural network (FFNN) is a sequence of interconnected *layers* $\{l_1, l_2, \dots, l_n\}$. When the number of layers
 123 is important, the term *Deep Neural Networks* is used. In an NN, each layer holds one or many nodes, called *neurons*.
 124 The first layer l_1 is called the *input layer*, the last one l_n is the *output layer* and the remaining layers $l_i : 2 \leq i \leq n - 1$
 125 are referred to as *hidden layers*. Likewise, the nodes in the hidden layers are called *hidden nodes*. Each hidden node
 126 is associated with a *bias* and an *activation function*. The nodes of a layer $l_i \in \{l_2, l_3, \dots, l_n\}$ are connected to the nodes
 127 of the previous layer via *weighted edges*. That is to say, a neuron of layer l_i receives data from layer l_{i-1} , calculates
 128 the weighted sum of this data and adds a bias. An activation function is then applied, and the result is forwarded to
 129 interconnected neurons of the next layer l_{i+1} (more details are given below). The propagation of data from the input
 130 layer to the output layer, passing through multiple hidden layers, is called "feed-forward propagation". An NN is built
 131 upon a training phase that aims to recognize and encode the underlying input-output relationship (correlation) of a
 132 data set. To evaluate an NN model, the accuracy, which is the rate of correct predictions, is calculated. Fig. 1 shows a
 133 neural network of 4 layers: an input layer of 3 inputs, two hidden layers of 4 and 3 nodes, respectively, and a 2-node
 134 output layer.
 135
 136
 137
 138
 139
 140

141 An NN model can indeed be seen as a function $\mathcal{N} : D_x \rightarrow D_y$, where D_x is the input domain and D_y is the output
 142 domain of the model. For image classification for example, D_x is a matrix of pixel values representing an image, D_y is
 143 the set of all possible classes of these images. As an NN model consists of a sequence of n layers, \mathcal{N} can be considered
 144 as a composition of a set of functions $\{f_1, f_2, \dots, f_n\}$ where $f_i, 1 \leq i \leq n$ is the corresponding function of layer l_i . This
 145 can be written, formally, as: $\mathcal{N}(x) = f_n(f_{n-1}(\dots(f_1(x))\dots))$, where f_1 is the identity function. In the following, we give
 146 some formal definitions pertaining to NN concepts and properties that will be used later on in this paper.
 147
 148

149 *Definition 2.1.* For a layer $l_i : i \in \{1 \dots n\}$, we define the set of neurons of l_i by S_i , with $|S_i|$ the number of neurons in
 150 the layer l_i . And for a neuron $n_{ij} \in S_i$, its value w.r.t to an input x is $v_{ij}(x)$. For simplicity, when x is not specific, we
 151 use v_{ij} instead of $v_{ij}(x)$.
 152
 153

154 Let $n_{ij} \in S_i$ be a neuron of a hidden layer l_i , its value v_{ij} is calculated in two steps:
 155
 156





Fig. 1. Example of a neural network

Fig. 2. An example showing the connection between a neuron of l_i and l_{i-1}

- (1) **Affine transformation:** calculates the sum of previous layer's outputs modulated by the weights assigned to the corresponding edges, plus the bias. This can be formulated as:

$$z_{ij} = \sum_{k=1}^{k=|S_{i-1}|} w_{j,k}^{i-1} \times v_{i-1,k} + b_{ij}$$

where $w_{j,k}^{i-1}$ is the weight of the edge connecting the nodes $n_{i-1,k}$ and n_{ij} , and b_{ij} is the bias of the node n_{ij} . Note that z_{ij} is also called the *pre-activation* value of n_{ij} .

- (2) **Activation function:** the final value v_{ij} , also called the value after activation, is determined by applying an activation function σ to z_{ij} , i.e. $v_{ij} = \sigma(z_{ij})$.

The two steps are summarized in Equation (1). The obtained value v_{ij} is the output value of n_{ij} which will be forwarded to the next layer l_{i+1} . Fig. 2 illustrates these steps on an example.

$$v_{ij} = \sigma \left(\sum_{k=1}^{k=|S_{i-1}|} w_{j,k}^{i-1} \times v_{i-1,k} + b_{ij} \right) \quad (1)$$

The calculation of the NN output $y = \mathcal{N}(x)$ for a given input x , is done by successively applying these operations, layer by layer, from the input to the output layer.

Depending on the application, there exists several activation functions: *Sigmoid*, *Tanh*, *Relu*, etc. [70]. *Relu* (for Rectified Linear Unit), as defined in Equation (2), is a piece-wise linear function that has linear behaviors in $(-\infty, 0]$ and in $[0, +\infty)$. The *Relu* activation function is widely used in NN, and due to its simple form and its piece-wise linear behaviour, the majority of the existing neural network verification and abstraction approaches consider models with this activation function [23, 30, 42].

$$\text{Relu}(x) = \max(x, 0) = \begin{cases} x, & \text{if } x \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Remark (Weights). In this paper, the weight of an edge connecting $n_{ik} \in S_i$ to a node $n_{i+1,j} \in S_{i+1}$ is written as w_{jk}^i or $w(n_{ik}, n_{i+1,j})$.

2.2 Verification of neural networks

Formal verification is the domain of proving or disproving that a system meets certain formal specifications and properties. A verification problem is defined as:

$$M \models P ? \quad (3)$$

which is equivalent to answering the question: does the system model M satisfy the property P ? Depending on the verification technique, the system has to be modelled (e.g., state transition model) and the specifications need to be expressed respecting some specific syntax (e.g., temporal logic). The aim of a verification technique is to prove that P holds on M or generate a counterexample witnessing the violation of P . Many verification techniques, such as model-checking, SAT/SMT, abstract interpretation, and theorem proving have been broadly and successfully applied to verify software-intensive systems [3, 10].

Accordingly, formal verification for NN can be defined as in Formula (3), where M is the NN model and P is the property to be checked, which is generally a mathematical formula constituted of a set of constraints on the inputs and the outputs of the network.

According to Leofante et al. [37], for a given NN represented by its corresponding function $\mathcal{N} : D_x \rightarrow D_y$, the NN verification problem can be stated as follows:

- Define $pre(x)$ and $post(y)$ as a set of constraints on the input x (preconditions) and the output y (postconditions), respectively. Here, $pre(x)$ and $post(y)$ are sorted first order logic formulas.
- For all x satisfying the preconditions $pre(x)$, verify whether or not $\mathcal{N}(x)$ fulfills the postconditions $post(\mathcal{N}(x))$.

This can be formulated as follows:

$$\forall x \in D_x, pre(x) \implies post(\mathcal{N}(x)) \quad (4)$$

Example 2.2. By taking $D_x = \mathbb{R}^2$ and $D_y = \mathbb{R}$ as the input and output domains of some given NN, the verification problem defined by Formula (4) can be instantiated as:

$$\begin{cases} pre(x) : x_1 \in [l_1, u_1] \wedge x_2 \in [l_2, u_2], \text{ with } x = (x_1 \ x_2)^T \\ post(\mathcal{N}(x)) : \mathcal{N}(x) \geq c \end{cases}$$




Fig. 3. An example of state-space explosion. For two *Relu* nodes, case splitting leads to four linear subproblems.

where $l_i, u_i, c \in \mathbb{R}$, and $l_i \leq u_i$. The verification problem of this example thus aims to check that for all input x in the 2-dimensional interval defined in the precondition, the corresponding output $\mathcal{N}(x)$ is lower-bounded by c as in the postcondition.

Example 2.3. To verify the robustness property of a classification network for an input image x_0 , i.e., to check for a classification problem that the network assigns the same label (class) c_i to all inputs within a small region surrounding x_0 , the verification problem can be formulated using (4) as follows:

$$\begin{cases} \text{Lets } x_0 \in D_x : \mathcal{N}(x_0) = c_i \\ \text{pre}(x) : \|x - x_0\|_p \leq \epsilon \\ \text{post}(\mathcal{N}(x)) : \mathcal{N}(x) = c_i \end{cases}$$

where $\|\cdot\|_p$ is a given norm.

It is worth mentioning that the paper [37] introduces other types of properties such as equivalence between two NNs. However, it should be noted that most of the existing verification methods and all the abstraction methods reviewed in this paper concentrate on verifying a single network and rely on properties based on Formula 3.

Verifying properties of NNs is increasingly receiving more attention and many approaches have been proposed in recent years [23, 40]. The straightforward verification way consists of encoding the NN behavior, as well as the property to be checked, as a system of linear equations, and then use an appropriate engine to perform the verification process. For instance, SAT/SMT and MILP encoding are widely used to verify NNs properties [8, 13, 24, 30, 31, 42, 58]. These methods are also called *complete* because they encode the exact behavior of the network. However, since most of the common activation functions are nonlinear, this kind of verification methods does not scale in the case of large neural networks, and suffers from state-space explosion. For example for the piece-wise linear activation function *Relu*, each *Relu* node has to be split into two linear constraints, i.e.: if $y = \text{relu}(x)$, then $y = 0$ when $x < 0$ and $y = x$ when x is positive. Therefore, solving a verification problem of a network of n *Relu* nodes leads to solving 2^n linear sub-problems as illustrated in Fig. 3. To address this issue, several approaches based on abstraction have been proposed. The next section provides more details about this category of techniques.




Fig. 4. The activation function Sigmoid (σ) and its abstraction in $x \in [-2, 2]$. The solid line represents $y = \sigma(x)$ and each small region (yellow rectangles) is an over-approximation of y [50].

3 ABSTRACTION APPROACHES FOR NEURAL NETWORK VERIFICATION

In order to overcome the drawbacks of complete verification methods for NN, some abstraction approaches are proposed. The main idea behind these approaches consists in generating an abstract model from the original network ensuring that whenever the property P holds on the abstract model \bar{N} , it necessarily holds on the original one N , i.e.,:

$$\bar{N} \models P \implies N \models P. \quad (5)$$

However, these approaches may fail to provide any conclusion on the original network when the property is violated on the abstract model. This is in fact due to spurious counterexamples. Namely, when the property does not hold, a counterexample (CE) on the abstract model is generated, but due to the over-approximation of the abstract model, this CE might not correspond to any real behavior in the original model (i.e., spurious counterexample).

Concretely, the abstraction of NN can be performed in two different manners:

- *Activation function abstraction* (AF abstraction): to ease the verification process, non-linear activation functions of the NN are over-approximated by a set of linear constraints.
- *NN model reduction*: abstracting the network model by merging some nodes in order to reduce the size of the network, and thus improve the scalability of existing verification tools.

A detailed survey of these methods is given in Sections 3.1 and 3.2, respectively.

Remark (Refinement). Some works consider improving the *incomplete* verification methods by ruling out as many spurious CE as possible by introducing a refinement phase. In other words, the verification method refines the abstract model iteratively until we can prove either the property holds or the generated CE exhibits a real behavior on the original model [15, 59, 60, 64, 65].

3.1 Abstraction of the activation function

The key challenge of NN verification is pertaining to the non-linearity of activation functions. AF abstraction-based verification approaches are applied to handle this issue by over-approximating the activation functions with linear constraints.

The earliest work dealing with NN verification problem was introduced by Pulina et al. [50]. In this work, authors divided the *Sigmoid* function into small regions, then a linear over-approximation is computed for each region, as shown in Fig. 4.

With the same spirit, Ehlers [15] proposed a precise *Relu*-abstraction technique where *Relu* is replaced by a system of linear constraints (see Fig. 5a) and hence the verification problem of NN is reformulated as a linear programming

(LP) problem that can be solved using classic LP solvers. The approach in [15] was implemented in a tool called Planet and brings the LP toolkit GLPK into play along with the Minisat solver for verification.

Gehr et al. [18] applied an *abstract interpretation* method [11] on NN for the first time. They proposed a framework called AI^2 (Abstract Interpretation for Artificial Intelligence) that soundly over-approximates NN operations by means of *zonotope* abstract domain³. The approach can be extended to support other abstract domains. AI^2 can handle feed-forward and convolutional neural networks (CNN) with *Relu* and *max-pooling* functions. The approach in [18] was extended by Singh et al. [54] to support *Sigmoid* and *Tanh* activation functions. This is accomplished by means of abstract transformers based on zonotopes for each function. As an example, the abstraction of *Relu* is given in Fig. 5b

Furthermore, Singh et al. [55] proposed a new method, called *DeepPoly*, based on Abstract Interpretation by introducing a new abstract domain. DeepPoly combines floating point polyhedra and intervals. Each neuron is represented by its concrete and symbolic upper and lower bounds. Moreover, Singh et al. [55] introduced abstract transformers for popular NN operations: affine transformation, *Relu*, *Sigmoid*, *Tanh* and *Max-pooling* to propagate the inputs successively through the layers of the network. For *Relu*, two different abstractions are proposed as shown in Figs. 5c and 5d. It is worthwhile to mention that the approach supports both feed-forward and convolutional NN.

While the previous works consider only a single neuron, some others try to define sound approximations of a set of neurons, jointly. Singh et al. [53] introduced a new method that provides an approximation of k *Relu* nodes (in the same layer) at a time in order to capture dependencies of the *Relu* inputs. First, the k nodes are selected and then the convex relaxation of the group of nodes is calculated. The framework has a parameter k which represents the number of *Relu* nodes to be considered together. A more general framework, based on [53], was recently proposed by Müller et al. [44]. The framework, called *PRIMA* (PReclse Multi-neuron Abstraction), computes the convex over-approximation of a set of k outputs of arbitrary activation function, including *Relu*, *Sigmoid* and *Tanh*. The approach decomposes the n activations into overlapping groups of size k , then calculates the convex approximation of the octahedral over-approximation for each group i . Finally, it takes the union of all the obtained output constraints. These constraints combined with the encoding of the whole NN are used for verification.

Other techniques based on symbolic propagation are proposed in [38, 69] to enhance the precision of abstract interpretation-based approaches. In symbolic propagation every neuron is associated with a formula expressed using the activations of neurons in its previous layers. In [56], a combination of over-approximation techniques with linear relaxation methods is proposed so as to gain more precision of over-approximation techniques and the scalability of complete methods.

These techniques can be adapted to support further types of NNs. One way to deal with Recurrent Neural Networks is to generate an equivalent feed-forward neural network and then apply the abstraction method [1, 27]. For Convolutional Neural Networks, most of the techniques are applicable and the only restriction is that the activation function of the convolution layer has to be *Relu* or other supported functions such as *Sigmoid* and *Tanh* [55].

3.2 NN model reduction

The main objective of NN model reduction is to reduce the size of the NN model while guaranteeing some behavioral relation: the desired property P holds on the original model N whenever it holds on the reduced model \bar{N} as defined in (5). Fig. 6 provides an illustrative example of the main idea behind model reduction applied on a small neural network.

³An abstract domain is a set of logical constraints that define a geometric shape. The most popular abstract domains are: box (or Interval), zonotope and polyhedra. For example, a zonotope abstract domain [19] Z is defined by a set of constraints z_i , s.t: $z_i = a_i + \sum_{j=1}^m b_{ij}\epsilon_j$, where $\epsilon_j \in [l_j, u_j]$ is an error term and a_i, b_{ij} are constants.

417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468

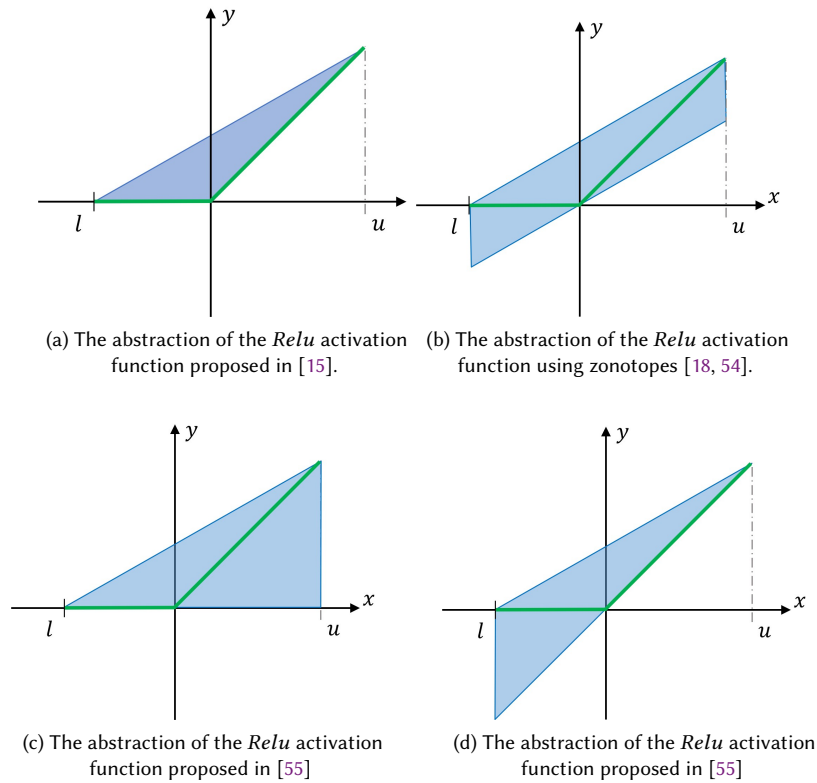


Fig. 5. *Relu* activation function abstractions using different abstract domains. The *Relu* ($y = \text{relu}(x)$) is represented by the green line and its over-approximation on the range $x \in [l, u]$ by the blue filled area.




Fig. 6. Model reduction of a small neural network

Such a behavioral relation is obtained by ensuring that \bar{N} is an over-approximation of N (i.e. all behaviors of N can be reproduced in \bar{N}). Therefore, the reduction process must carefully select the set of neurons to be merged (or removed), and determine how to calculate the weights of the new edges.

Prabhakar and Afzal [49] proposed a method based on Interval Neural Networks (INN) for output range analysis. In this method, the nodes of the same layer are merged while replacing the weights of their input edges by the interval hull

469 of the incoming edges. In other words, the weights of incoming edges are replaced by $[\min(W_{in}), \max(W_{in})]$, where
 470 W_{in} are the values of the incoming weights to the nodes to be merged. The weights of the outgoing edges from these
 471 nodes are replaced by the interval hull multiplied by the number of merged nodes $n: n \times [\min(W_{out}), \max(W_{out})]$.
 472

473 For the verification part, Prabhakar and Afzal [49] adapted INN to MILP big-M encoding [8] and used the Gurobi
 474 MILP solver for verification. The performance of this method is tested on the airborne collision avoidance ACAS Xu
 475 benchmark [29, 30]. The authors claim that the abstraction enhances the verification process. Namely, Gurobi was not
 476 able to verify a number of properties on the original model (no return), while the same properties have been successfully
 477 checked when Gurobi was applied on the abstract model.
 478

479 Recently, Boudardara et al. [6] proposed an interval-weight based model reduction method. The elaborated method
 480 supports *Relu* and *Tanh* neural networks. While an outgoing weight of a set of merged nodes is the sum of absolute
 481 values of their corresponding outgoing weights, an abstract incoming weight is an interval defined as: the min and the
 482 max of the sign of the corresponding outgoing weights of the merged nodes multiplied by the the original incoming
 483 weights. The sign function defined in this work returns 1 if the value is at least equal to 0, and -1 otherwise. The method
 484 is applied to the ACAS Xu *Relu*-NN benchmark [30], where the Interval Bound Propagation (IBP) algorithm [67] is
 485 used to calculate the output range on the original and abstract networks. Moreover, the authors of [6] have varied the
 486 number of merged nodes to assess the output range and the IBP computation time for abstract networks. It has been
 487 shown, in particular, that merging more nodes accelerates the IBP algorithm while generating larger output ranges.
 488 The authors observe that the wide output range is due to the IBP algorithm which is not an exact verification method.
 489 However, this work does not discuss possible adaptation to support other verification tools.
 490

491 In [57], Sotoudeh and Thakur, by introducing the notion of Abstract Neural Network (ANN), provided a formalization
 492 of a general abstraction approach. In ANN, the weights are represented using abstract domains. Accordingly, the
 493 approach proposed by Prabhakar and Afzal [49] can be considered as a particular instantiation of this approach using
 494 the interval abstract domain. Notice that the proposed approach supports a wide range of activation functions. Moreover,
 495 it can be instantiated using other convex abstract domains and it is not restricted to intervals as used in INNs [49]. The
 496 approach provides a generic formula to calculate the weight merging matrix \overline{W} from the original weight matrix W and
 497 the partitions P^{in} and P^{out} of two successive abstract layers l_i and l_{i+1} , respectively. A partition P_i is a rearrangement
 498 of a set S_i of neurons, i.e., if $S_i = \{n_{i1}, n_{i2}, n_{i3}\}$, a possible partition of S_i would be $P_i = \{\{n_{i1}, n_{i2}\}, \{n_{i3}\}\}$, which means
 499 that n_{i1} and n_{i2} will be merged in the abstract network. \overline{W} is the convex combination (calculated by a function g) of
 500 the partitioning combination matrix of P^{in} and P^{out} , denoted by C and D , respectively, and the weight matrix W , i.e.,
 501 $\overline{W} = g(D, W, C)$. Next, the abstract weight matrix, denoted by W_{abs} , is built by applying a convex abstract domain α_A
 502 on the obtained \overline{W} : $W_{abs} = \alpha_A(\overline{W})$. The reduced model is obtained by applying the same procedure to every layer,
 503 iteratively. Therefore, the obtained reduced model is an over-approximation for any non-negative activation function
 504 that satisfies the Weakened Intermediate Value Property (WIVP). Although some activation functions can have negative
 505 values and others are not continuous (thus not WIVP), the authors of [57] claim that there is always a way to overcome
 506 these problems, as they showed for Leaky *Relu* and the *threshold* activation functions.
 507

508 In [2], Ashok et al. apply K-means clustering algorithms to partition each hidden layer l_i into k_i subgroups, such that
 509 $k_i \leq |S_i|$, then replace each subgroup with its representative neuron. The abstraction method, called DeepAbstract, has
 510 three parameters: the original network N , a finite set of input-points X and a vector K_L which contains the number of
 511 nodes on each abstract layer. For each hidden layer l_i , the following steps are performed:
 512

- 513 (1) For every $x \in X$, calculate the value $v_{ij}(x)$ of each neuron in S_i ,

- 521 (2) Apply K-means to split each layer l_i into k_i clusters. Let C_{l_i} denote the set of clusters of l_i ,
 522 (3) For each cluster $C \in C_{l_i}$:
 523 (a) Determine the representative neuron rep_C ,
 524 (b) Calculate the corresponding outgoing weights of rep_C :

$$525 \quad \bar{W}_{*, rep_C}^i = \sum_{n_{ij} \in C} W_{*, n_{ij}}^i$$

- 526
 527
 528
 529 (c) Replace all the neurons in C with rep_C .

530 Note that the representative neuron rep_C of a cluster C is the nearest neuron to the centroid of C , thus; the incoming
 531 weights of rep_C remain the same as the corresponding neuron before abstraction. All the other neurons from cluster C
 532 are omitted with their incoming edges.
 533

534 In addition, Ashok et al. [2] provide a method to lift the verification results from the abstract model to the original
 535 one. The idea is to calculate the accumulated error induced by replacing a cluster of neurons by its representative
 536 for each image x in X , and then propagate this error through the successive layers using the DeepPoly verification
 537 Algorithm⁴. A set of experiments were conducted to check the performance of DeepAbstract. Local robustness of some
 538 MNIST images was checked and the authors claim that the verification time was reduced by 25% when DeepPoly is
 539 combined with DeepAbstract.
 540

541 Elboher et al. [16] proposed an abstraction approach based on merging neurons of the same *category* (see hereafter)
 542 to build a smaller model so as to enhance the scalability of the existing verification tools. Regarding the verification
 543 property, which has the form: $P : \forall x \in pre(x) \implies y \leq c$, the aim of this approach is to build a reduced model \bar{N}
 544 (its corresponding function is \bar{N}), s.t $\forall x \in D_x, \bar{N}(x) \geq N(x)$. Therefore, $N \models P$ whenever $\bar{N} \models P$ (i.e., $\bar{N}(x) \leq c$).
 545 First, each neuron is labelled according to the sign of its outgoing weights. A neuron is split if it has both positive and
 546 negative outgoing weights. Next, to guarantee that \bar{N} is an over-approximation of N , the proposed method tries to
 547 increase the output of the abstract model by classifying each neuron as I or D . The class I means the output will increase
 548 by increasing the value of this neuron, while a neuron is marked as D when decreasing its value leads to increasing the
 549 output's value. Finally, the nodes of the same layer and the same category can be merged by summing up the weights of
 550 their outgoing edges and taking the *min* value of the the weights of their incoming edges if they are marked as D , or
 551 the *max* value for any I group of nodes. Moreover, some heuristics are proposed in [16] to enhance the abstraction
 552 process. The proposed method is applied on ACAS Xu networks while Marabou [31] is used as back-end verification
 553 tool. A comparison study between the abstraction method combined with Marabou and the vanilla version of Marabou
 554 was conducted, and the results showed that the abstraction method allows Marabou to verify more properties in less
 555 execution time.
 556

557 A novel approach based on bisimulation [33] is proposed by Prabhakar [48]. The generated abstract neural network
 558 is equivalent, or bisimilar, to the original one. To guarantee the equivalence between N and \bar{N} , two neurons n_{ij} and
 559 n_{ik} to be merged must have the same activation function, the same bias value ($b_{ij} = b_{ik}$) and the same weights for
 560 each incoming edge respectively, i.e., $\forall n' \in S_{i-1}, w(n', n_{ij}) = w(n', n_{ik})$. Due to the strict conditions that, generally,
 561 do not hold in most of real networks, Prabhakar [48] extends the NN bisimulation to a more feasible relaxed method,
 562 called NN δ -bisimulation. Using NN δ -bisimulation ($\delta \in \mathbb{R}^+$), two nodes n_{ij} and n_{ik} in S_i can be merged if the following
 563 conditions are satisfied:
 564

- 565 (1) n_{ij} and n_{ik} have the same activation function
 566

567
 568
 569
 570
 571 ⁴Available at <https://github.com/eth-sri/ERAN>.

Table 1. A list of NN model reduction methods used for verification. The underscore symbol “-” is used to denote that no information is provided in the corresponding original paper.

Method	Pub. Year	Supported AFs	Verification methods	Evaluation on	Guarantees of the reduced model
R4V [52]	2019	<i>Relu</i>	Relupex[30], ERAN[54], Neurify[64], Planet[15]	DAVE-2[4], DroNet[43]	None
INN [49]	2019	<i>Relu</i>	MILP [42]	ACAS Xu [30]	$\mathcal{N}(x) \in \overline{\mathcal{N}}(x)$
ANN [57]	2020	<i>Relu</i> , <i>Leaky Relu</i> ⁵	-	-	$\mathcal{N}(x) \in \overline{\mathcal{N}}(x)$
DeepAbstract [2]	2020	<i>Relu</i>	ERAN	MNIST[35]	Depends on the data set
Elboher et al. [16]	2020	<i>Relu</i>	Marabou[31]	ACAS Xu [30]	$\mathcal{N}(x) \leq \overline{\mathcal{N}}(x)$
Bisimulation [48]	2021	<i>Relu</i>	-	-	$\mathcal{N} \equiv \overline{\mathcal{N}}$ ⁶
Boudardara et al. [6]	2022	<i>Relu</i> , <i>Tanh</i>	IBP[67]	ACAS Xu [30]	$\mathcal{N}(x) \in \overline{\mathcal{N}}(x)$

$$(2) |b_{ij} - b_{ik}| \leq \delta$$

$$(3) \forall n' \in S_{i-1}, |w(n', n_{ij}) - w(n', n_{ik})| \leq \delta$$

where $\delta \geq 0$. So the obtained network \overline{N} is δ -bisimilar to network N .

Taking advantages of code refactoring [17], Shriver et al. [52] introduced the concept of refactoring neural networks to restructure the initial model and preserve its accuracy to enhance further operations on it, for instance verification. Concretely, NN refactoring consists of two steps: architecture transformation and distillation. The former applies some changes on the network’s architecture by dropping or changing some layers and/or their types that are not supported by verification tools (e.g. residual blocks and convolutional layers). The latter updates the model’s parameters: weights and biases, while preserving the original model’s behavior, which is captured by its accuracy and test error according to Shriver et al. [52]. A tool called R4V was developed from this approach. R4V was tested on DAVE-2 [4] and DroNet [43] networks. The used verification tools are presented in Table 1. The results showed that applying the verification tools on the refactored model improves their scalability. For example, Planet [15] fails to check any property on DroNet within 24 hours. However, after refactoring the network, Planet was able to verify three out of the ten properties.

The main features of the above discussed neural networks reduction techniques are summarized in Table 1. The last two columns of the table contain verification methods and the data sets used during the evaluation of the abstraction method. Verification methods are those used during the evaluation of the abstraction in the original paper; notice that other methods can be used to verify the obtained abstract model.

An example is provided in Figure 7 to demonstrate the application of some of the methods mentioned in this section [6, 16, 49, 57]⁷. Notice that the abstract network using the ANN method [57] with the box (or interval) abstract domain is the same as the abstract network obtained using the method of INN [49] (see Figure 7b). The example presents a segment of a *Relu*-NN, i.e. s_1 is an arbitrary neuron of a hidden layer and not the input of the network, and all nodes are assigned a *Relu* activation function. We apply abstraction (using the selected methods) to merge the two nodes s_3

⁵The authors claim that the method can be adjusted to support other activation function

⁶The abstract network is equivalent to the original one when bisimulation is used which is not the case for δ -bisimulation

⁷Supplementary details are needed to apply the other methods. For example DeepAbstract [2] needs a data set for the clustering algorithm.

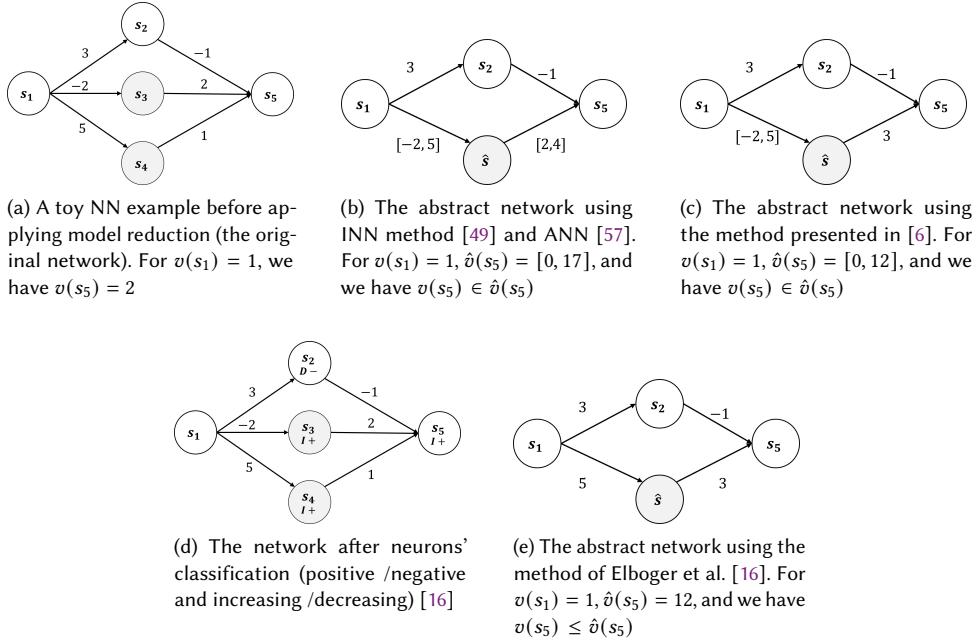


Fig. 7. The application of different model reduction methods on a toy example of NN

and s_4 , while assuming that $v(s_1) = 1$, and we calculate the value of s_5 , $v(s_5)$ and $\hat{v}(s_5)$ on the original and the abstract networks, respectively. While model reduction methods [6, 49, 57] (Figures 7c and 7b) ensure that the output of the original network is within the ranges of the output of the abstract network, i.e.: $v(s_5) \in \hat{v}(s_5)$, the method introduced in [16] (Figure 7e) guarantees that the output of the obtained abstract network is always greater than the output's value of the original network, i.e.: $v(s_5) \leq \hat{v}(s_5)$

It is worth noting here that these techniques can be adjusted to support other types of NNs. For instance, RNN can be transformed into an equivalent FFNN [1, 27], and then model reduction approaches can be applied to generate the abstract network. On the other hand, model reduction can be applied on the fully connected part of CNNs [46, 68]. Regarding Binarized Neural Networks (BNN), due to their binary behaviour and their small size comparing to other types of NNs, their verification does not require abstracting their behavior and, generally, exact methods such as SAT and MILP can be applied directly [28, 34, 45].

It is worth mentioning that another family of techniques based on merging neurons and removing some edges without affecting the accuracy of the model exists in the literature. These techniques are called NN compression and acceleration, and their objective is to build a smaller network with low computational complexity, so that it can be embedded on devices with limited resources and used in real-time applications, while keeping the accuracy as high as possible [9, 21, 39]. Although both NN model reduction and NN compression strive to reduce the number of neurons, NN compression techniques cannot be used for verification, since the generated models do not fulfil the abstraction condition presented in Formula (5). In other words, verifying a property P on the compressed network obtained by any compression method does not imply that the property does hold in the original network.

3.3 Discussion

This section discusses the aforementioned model reduction methods, while highlighting their limitations and proposing some possible area of improvements. In order to fairly compare the efficiency of the discussed approaches, we analyze them according to three main criteria (with respect to the available information in the original papers): (i) the precision of the over-approximation, (ii) the minimal number of neurons that can be obtained when the reduction method is applied until saturation, and (iii) the efficiency regarding the verification time and the number of the verified properties on the reduced model versus the original one.

The abstraction method based on INNs, proposed by Prabhakar et al. [49] seems to be very useful when the problem of output range analysis is considered. An exhaustive application of this method leads to merge all neurons of each hidden layer and replace them by one abstract neuron. The results of their paper show that the precision depends highly on the number and the selection of the nodes to be merged. The method needs some improvements to be more precise, since no study has been provided for neuron selection. In addition, operations on intervals may impact the precision of this method. MILP encoding is proposed to solve the verification problem on INNs, and to the best of our knowledge, no other verification method is proposed to verify INNs. Moreover, this method is restricted to abstract NNs with non-negative activation functions [57]. Consequently, Sotoudeh et al. [57] proposed some fundamentals to abstract any NNs with different activation functions using any convex abstract domain and which is not limited to intervals. In [57], the authors provide an example of abstraction based on octagons, but no explanation was given of the meaning of using such abstract domain to represent the merged neurons. Moreover, the work would have been more relevant if it had included an evaluation study to concretely show how the ANN can be extended to deal with other abstract domains. In [6], Boudardara et al. proposed a method that is similar to INNs[49], where the incoming weights are encoded as intervals, while the outgoing weights are scalars. However, unlike INNs, the proposed method is not limited to non-negative activation functions and can support the use of *Tanh* activation function. Moreover, the authors claim that the method can be adjusted to support other activation functions as well.

DeepAbstract, proposed by Ashok et al. [2], is parametrized by the number of clusters on each layer; if there are few clusters, the model will be more abstract and less precise. In addition, this method relies on the discrete input set X that is used during clustering phase and can only verify the robustness of the model on points within this set X . Ashok et al. [2] claim that the verification time was reduced by 25% when DeepAbstract is used along with DeepPoly, however, only 195 out of 200 images could be verified to be robust against 197/200 when DeepPoly is used without abstraction.

The abstraction-refinement proposed by Elboher et al. [16] boosted the Marabou verifier to check more properties (58 out of 90 property versus 35/90). Moreover, the abstraction method reduces the total query median runtime from 63671 seconds to 1045 seconds. As a consequence of the classification of neurons, this method can abstract a layer to four neurons at most. This is one of the main drawbacks of this method since only neurons belonging to the same category can be merged. It should also be mentioned that only properties in the form: $y \leq c$ are considered, although authors claim that the approach is adaptable to cope with various types of properties by adjusting the output layer. In addition, this method cannot be applied if some neurons have negative values. For instance, this method cannot be applied in hidden layers if the used activation function returns negative values such as sigmoid and Leaky *Relu*. For the same reason, the first hidden layer cannot be abstracted if the inputs are negative. An example demonstrating this case is given in Fig. 8, where x is an input, y is the output.

The NN in Fig. 8.b is generated using Elboher et al.'s method [16], which is supposed to be an abstraction of the original model of Fig. 8.a. Both \bar{N} and N use the *Relu* activation function on the hidden layer. Although for negative




Fig. 8. Counterexample of Elboher et al. [16] abstraction method

inputs the output of \bar{N} is always zero: $\forall x \leq 0, \bar{y} = 0$, the output of N is always positive, for instance, for $x = -1, y = 3$, thus the condition of the over-approximation $\forall x \in D_x : \bar{N}(x) \geq N(x)$ does not hold.

The NN bisimulation method proposed in [48] guarantees the equivalence between abstract and original models, thus offers an exact abstraction. However, the set of conditions are hard to satisfy on real neural network, especially the condition on weights. On the other hand, the relaxed version, NN- δ -bisimulation, looks more feasible but needs further improvements to keep trace of the verified property on the abstract model and lift it to provide guarantees on the original network.

In [52], Shriver et al. propose an efficient approach with a dedicated tool, called R4V, to simplify and compress NN models. The wide experimental study they performed with different verification tools and data sets shows that R4V offers actual benefits to overcome the limitations of some NN verification techniques. However, this method enables to verify properties on the refactored model and does not propose a way to lift these guarantees to the original model. In other words, it does not provide any guarantee of whether the property holds on the original model.

Regarding the challenges of neural network verification, developing a new general approach that overcomes the issues related to the existing abstraction methods mentioned above is necessary. The works [2, 48, 52] could be adopted and combined with some heuristics to select candidate neurons to be merged. For instance, the δ -bisimulation method [48] can be used to select similar nodes by analyzing their incoming weights. The approach in [2] can be adapted using discretization of the input region, so that the nodes that are close to each other (in the same cluster) are good candidates for abstraction.

While the technique in [16] ensures that $N(x) \leq \bar{N}(x)$, the three methods presented in [6, 49, 57] go further by guaranteeing that the output of the original network is always included within the output range of the obtained abstract network, i.e., $N(x) \in \bar{N}(x)$. However, it is necessary to conduct a comparative study to assess the performance of these methods. On the other hand, an abstract network obtained using DeepAbstract [2] can be used only to verify the robustness of the model on inputs within the set of images X that is used during the clustering phase. The last column of Table 1 summarizes the relation between the original and the abstract networks using different methods.

4 CONCLUSION

In this work, we discussed the problem of neural network verification and we presented different existing techniques used to solve this problem. We showed that the abstraction of neural networks can be used to help tackle the non-linearity and the complexity of the generated models. Abstraction of neural networks can be applied in two levels: abstracting

781 the activation function and reducing the network’s size (model reduction). While the abstraction of activation functions
 782 aims to over-approximate the non-linear activation functions with linear constraints, model reduction is used to reduce
 783 the number of neurons of the network. Both categories are applied to improve the verification process as a whole. The
 784 abstraction has to be sound, meaning that the desired behavior of the original model must be maintained. In this paper
 785 we focused more on model reduction methods since, to the best of our knowledge, no survey about neural networks
 786 reduction for verification purposes has been introduced.
 787

788 While the main focus of this work is on the application of abstraction methods to feed-forward neural networks,
 789 discussing their advantages, limitations, and the formal guarantees provided by each model reduction method, we also
 790 addressed the perspectives and potential applicability of these methods to other types of NNs, including convolutional
 791 neural networks (CNNs) and recurrent neural networks (RNNs).
 792
 793

794 ACKNOWLEDGMENTS

795 This research work is funded by the French program “Investissements d’Avenir” and is part of the French collaborative
 796 project TASV (Train Autonome Service Voyageurs), with SNCF, Alstom Crespín, Thales, Bosch, and Spiros
 797
 798

800 REFERENCES

- 801 [1] Michael E Akintunde, Andreea Kevorchian, Alessio Lomuscio, and Edoardo Pirovano. 2019. Verification of RNN-based neural agent-environment
 802 systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 6006–6013.
- 803 [2] Pranav Ashok, Vahid Hashemi, Jan Křetínský, and Stefanie Mohr. 2020. Deepabstract: Neural network abstraction for accelerating verification. In
 804 *International Symposium on Automated Technology for Verification and Analysis*. Springer, 92–107.
- 805 [3] Armin Biere, Marijn Heule, and Hans van Maaren. 2009. *Handbook of satisfiability*. Vol. 185. IOS press.
- 806 [4] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prason Goyal, Lawrence D Jackel, Mathew Monfort, Urs
 807 Muller, Jiakai Zhang, et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316* (2016).
- 808 [5] Benedikt Bollig, Martin Leucker, and Daniel Neider. 2022. A Survey of Model Learning Techniques for Recurrent Neural Networks. *A Journey from*
 809 *Process Algebra via Timed Automata to Model Learning: Essays Dedicated to Frits Vaandrager on the Occasion of His 60th Birthday (2022)*, 81–97.
- 810 [6] Fateh Boudardara, Abderraouf Boussif, Pierre-Jean Meyer, and Mohamed Ghazel. 2022. Interval Weight-Based Abstraction for Neural Network
 811 Verification. In *Computer Safety, Reliability, and Security. SAFECOMP 2022 Workshops: DECSoS, DepDevOps, SASSUR, SENSEI, USDAI, and WAISE*
 812 *Munich, Germany, September 6–9, 2022, Proceedings*. Springer, 330–342.
- 813 [7] Xuyi Cai, Ying Wang, and Lei Zhang. 2022. Optimus: An Operator Fusion Framework for Deep Neural Networks. *ACM Transactions on Embedded*
 814 *Computing Systems (TECS)* (feb 2022). <https://doi.org/10.1145/3520142>
- 815 [8] Chih-Hong Cheng, Georg Nührenberg, and Harald Ruess. 2017. Maximum resilience of artificial neural networks. In *International Symposium on*
 816 *Automated Technology for Verification and Analysis*. Springer, 251–268.
- 817 [9] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. 2017. A survey of model compression and acceleration for deep neural networks. *arXiv preprint*
 818 *arXiv:1710.09282* (2017).
- 819 [10] Edmund M Clarke, Thomas A Henzinger, Helmut Veith, Roderick Bloem, et al. 2018. *Handbook of model checking*. Vol. 10. Springer.
- 820 [11] Patrick Cousot and Radhia Cousot. 1977. Abstract interpretation: a unified lattice model for static analysis of programs by construction or
 821 approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*. 238–252.
- 822 [12] Souradeep Dutta, Susmit Jha, Sriram Sankaranarayanan, and Ashish Tiwari. 2017. Output range analysis for deep neural networks. *arXiv preprint*
 823 *arXiv:1709.09130* (2017).
- 824 [13] Souradeep Dutta, Susmit Jha, Sriram Sankaranarayanan, and Ashish Tiwari. 2018. Output Range Analysis for Deep Feedforward Neural Networks.
 825 In *Proc. 10th NASA Formal Methods*. 121–138. https://doi.org/10.1007/978-3-319-77935-5_9
- 826 [14] Krishnamurthy Dvijotham, Robert Stanforth, Sven Gowal, Timothy A Mann, and Pushmeet Kohli. 2018. A Dual Approach to Scalable Verification
 827 of Deep Networks.. In *UAI*, Vol. 1. 3.
- 828 [15] Ruediger Ehlers. 2017. Formal verification of piece-wise linear feed-forward neural networks. In *International Symposium on Automated Technology*
 829 *for Verification and Analysis*. Springer, 269–286.
- 830 [16] Yizhak Yisrael Elboher, Justin Gottschlich, and Guy Katz. 2020. An abstraction-based framework for neural network verification. In *International*
 831 *Conference on Computer Aided Verification*. Springer, 43–65.
- 832 [17] Martin Fowler. 2018. *Refactoring: improving the design of existing code*. Addison-Wesley Professional.
- 833 [18] Timon Gehr, Matthew Mirman, Dana Drachler-Cohen, Petar Tsankov, Swarat Chaudhuri, and Martin Vechev. 2018. Ai2: Safety and robustness
 certification of neural networks with abstract interpretation. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3–18.

- [19] Khalil Ghorbal, Eric Goubault, and Sylvie Putot. 2009. The zonotope abstract domain taylor1+. In *International Conference on Computer Aided Verification*. Springer, 627–633.
- [20] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [21] Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149* (2015).
- [22] Chao Huang, Jiameng Fan, Wenchao Li, Xin Chen, and Qi Zhu. 2019. ReachNN: Reachability Analysis of Neural-Network Controlled Systems. *ACM Transactions on Embedded Computing Systems (TECS)* 18, 5s, Article 106 (oct 2019), 22 pages. <https://doi.org/10.1145/3358228>
- [23] Xiaowei Huang, Daniel Kroening, Wenjie Ruan, James Sharp, Youcheng Sun, Emese Thamo, Min Wu, and Xinpeng Yi. 2020. A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability. *Computer Science Review* 37 (2020), 100270. <https://doi.org/10.1016/j.cosrev.2020.100270>
- [24] Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. 2017. Safety verification of deep neural networks. In *International conference on computer aided verification*. Springer, 3–29.
- [25] Radoslav Ivanov, Taylor J Carpenter, James Weimer, Rajeev Alur, George J Pappas, and Insup Lee. 2020. Verifying the safety of autonomous systems with neural network controllers. *ACM Transactions on Embedded Computing Systems (TECS)* 20, 1 (2020), 1–26.
- [26] Radoslav Ivanov, Kishor Jothimurugan, Steve Hsu, Shaan Vaidya, Rajeev Alur, and Osbert Bastani. 2021. Compositional Learning and Verification of Neural Network Controllers. *ACM Transactions on Embedded Computing Systems (TECS)* 20, 5s (2021), 1–26.
- [27] Yuval Jacoby, Clark Barrett, and Guy Katz. 2020. Verifying recurrent neural networks using invariant inference. In *Automated Technology for Verification and Analysis: 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19–23, 2020, Proceedings 18*. Springer, 57–74.
- [28] Kai Jia and Martin Rinard. 2020. Efficient exact verification of binarized neural networks. *Advances in neural information processing systems* 33 (2020), 1782–1795.
- [29] Kyle D Julian, Jessica Lopez, Jeffrey S Brush, Michael P Owen, and Mykel J Kochenderfer. 2016. Policy compression for aircraft collision avoidance systems. In *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*. IEEE, 1–10.
- [30] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. 2017. Reluplex: An efficient SMT solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*. Springer, 97–117.
- [31] Guy Katz, Derek A Huang, Duligur Ibeling, Kyle Julian, Christopher Lazarus, Rachel Lim, Parth Shah, Shantanu Thakoor, Haoze Wu, Aleksandar Zeljić, et al. 2019. The marabou framework for verification and analysis of deep neural networks. In *International Conference on Computer Aided Verification*. Springer, 443–452.
- [32] Jianglin Lan, Yang Zheng, and Alessio Lomuscio. 2022. Tight neural network verification via semidefinite relaxations and linear reformulations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 7272–7280.
- [33] Kim G Larsen and Arne Skou. 1991. Bisimulation through probabilistic testing. *Information and computation* 94, 1 (1991), 1–28.
- [34] Christopher Lazarus and Mykel J Kochenderfer. 2022. A mixed integer programming approach for verifying properties of binarized neural networks. *arXiv preprint arXiv:2203.07078* (2022).
- [35] Yann LeCun. 1998. *The MNIST database of handwritten digits*.
- [36] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [37] Francesco Leofante, Nina Narodytska, Luca Pulina, and Armando Tacchella. 2018. Automated verification of neural networks: Advances, challenges and perspectives. *arXiv preprint arXiv:1805.09938* (2018).
- [38] Jianlin Li, Jiangchao Liu, Pengfei Yang, Liqian Chen, Xiaowei Huang, and Lijun Zhang. 2019. Analyzing deep neural networks with symbolic propagation: Towards higher precision and faster verification. In *International Static Analysis Symposium*. Springer, 296–319.
- [39] Tailin Liang, John Glossner, Lei Wang, Shaobo Shi, and Xiaotong Zhang. 2021. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing* 461 (2021), 370–403.
- [40] Changliu Liu, Tomer Arnon, Christopher Lazarus, Christopher Strong, Clark Barrett, Mykel J Kochenderfer, et al. 2021. Algorithms for verifying deep neural networks. *Foundations and Trends® in Optimization* 4, 3-4 (2021), 244–404.
- [41] Weibo Liu, Zidong Wang, Xiaohui Liu, Nianyin Zeng, Yurong Liu, and Fuad E Alsaadi. 2017. A survey of deep neural network architectures and their applications. *Neurocomputing* 234 (2017), 11–26.
- [42] Alessio Lomuscio and Lalit Maganti. 2017. An approach to reachability analysis for feed-forward relu neural networks. *arXiv preprint arXiv:1706.07351* (2017).
- [43] Antonio Loquercio, Ana I Maqueda, Carlos R Del-Blanco, and Davide Scaramuzza. 2018. Dronet: Learning to fly by driving. *IEEE Robotics and Automation Letters* 3, 2 (2018), 1088–1095.
- [44] Mark Niklas Müller, Gleb Makarchuk, Gagandeep Singh, Markus Püschel, and Martin Vechev. 2022. PRIMA: general and precise neural network certification via scalable convex hull approximations. *Proceedings of the ACM on Programming Languages* 6, POPL (2022), 1–33.
- [45] Nina Narodytska, Shiva Kasiviswanathan, Leonid Ryzhyk, Mooly Sagiv, and Toby Walsh. 2018. Verifying properties of binarized deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [46] Matan Ostrovsky, Clark Barrett, and Guy Katz. 2022. An abstraction-refinement approach to verifying convolutional neural networks. In *Automated Technology for Verification and Analysis: 20th International Symposium, ATVA 2022, Virtual Event, October 25–28, 2022, Proceedings*. Springer, 391–396.
- [47] Ajeet Ram Pathak, Manjusha Pandey, and Siddharth Rautaray. 2018. Application of deep learning for object detection. *Procedia computer science* 132 (2018), 1706–1717.

- 885 [48] Pavithra Prabhakar. 2022. Bisimulations for neural network reduction. In *Verification, Model Checking, and Abstract Interpretation: 23rd International*
886 *Conference, VMCAI 2022, Philadelphia, PA, USA, January 16–18, 2022, Proceedings*. Springer, 285–300.
- 887 [49] Pavithra Prabhakar and Zahra Rahimi Afzal. 2019. Abstraction based output range analysis for neural networks. *Advances in Neural Information*
888 *Processing Systems* 32 (2019).
- 889 [50] Luca Pulina and Armando Tacchella. 2010. An abstraction-refinement approach to verification of artificial neural networks. In *International*
890 *Conference on Computer Aided Verification*. Springer, 243–257.
- 891 [51] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. 2018. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344* (2018).
- 892 [52] David Shriver, Dong Xu, Sebastian Elbaum, and Matthew B Dwyer. 2019. Refactoring neural networks for verification. *arXiv preprint arXiv:1908.08026*
(2019).
- 893 [53] Gagandeep Singh, Rupanshu Ganvir, Markus Püschel, and Martin Vechev. 2019. Beyond the Single Neuron Convex Barrier for Neural Network
894 Certification. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett
895 (Eds.), Vol. 32. Curran Associates, Inc.
- 896 [54] Gagandeep Singh, Timon Gehr, Matthew Mirman, Markus Püschel, and Martin Vechev. 2018. Fast and effective robustness certification. *Advances in*
897 *neural information processing systems* 31 (2018).
- 898 [55] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. 2019. An abstract domain for certifying neural networks. *Proceedings of the*
899 *ACM on Programming Languages* 3, POPL (2019), 1–30.
- 900 [56] Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. 2019. Boosting robustness certification of neural networks. In *International*
Conference on Learning Representations.
- 901 [57] Matthew Sotoudeh and Aditya V Thakur. 2020. Abstract Neural Networks. In *International Static Analysis Symposium*. Springer, 65–88.
- 902 [58] Vincent Tjeng, Kai Y. Xiao, and Russ Tedrake. 2019. Evaluating Robustness of Neural Networks with Mixed Integer Programming. In *International*
903 *Conference on Learning Representations*. <https://openreview.net/forum?id=HyGldiRqtm>
- 904 [59] Hoang-Dung Tran, Stanley Bak, Weiming Xiang, and Taylor T Johnson. 2020. Verification of deep convolutional neural networks using imagestars.
905 In *International Conference on Computer Aided Verification*. Springer, 18–42.
- 906 [60] Hoang-Dung Tran, Diago Manzananas Lopez, Patrick Musau, Xiaodong Yang, Luan Viet Nguyen, Weiming Xiang, and Taylor T Johnson. 2019.
907 Star-based reachability analysis of deep neural networks. In *International Symposium on Formal Methods*. Springer, 670–686.
- 908 [61] Hoang-Dung Tran, Weiming Xiang, and Taylor T Johnson. 2020. Verification approaches for learning-enabled autonomous cyber-physical systems.
909 *IEEE Design & Test* (2020).
- 910 [62] Damien Trentesaux, Rudy Dahyot, Abel Ouedraogo, Diego Arenas, Sébastien Lefebvre, Walter Schön, Benjamin Lussier, and Hugues Cheritel. 2018.
911 The autonomous train. In *2018 13th Annual Conference on System of Systems Engineering (SoSE)*. IEEE, 514–520.
- 912 [63] Caterina Urban and Antoine Miné. 2021. A Review of Formal Methods applied to Machine Learning. *arXiv preprint arXiv:2104.02466* (2021).
- 913 [64] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. 2018. Efficient formal safety analysis of neural networks. *arXiv preprint*
914 *arXiv:1809.08098* (2018).
- 915 [65] Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. 2018. Formal security analysis of neural networks using symbolic
916 intervals. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 1599–1614.
- 917 [66] Eric Wong and Zico Kolter. 2018. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International*
918 *Conference on Machine Learning*. PMLR, 5286–5295.
- 919 [67] Weiming Xiang, Hoang-Dung Tran, Xiaodong Yang, and Taylor T Johnson. 2020. Reachable set estimation for neural network control systems: A
920 simulation-guided approach. *IEEE Transactions on Neural Networks and Learning Systems* 32, 5 (2020), 1821–1830.
- 921 [68] Jin Xu, Zishan Li, Miaomiao Zhang, and Bowen Du. 2021. Conv-Reluplex: a verification framework for convolution neural networks. In *Proceedings*
922 *of the 33rd International Conference on Software Engineering and Knowledge Engineering (SEKE)*.
- 923 [69] Pengfei Yang, Jianlin Li, Jiangchao Liu, Cheng-Chao Huang, Renjue Li, Liqian Chen, Xiaowei Huang, and Lijun Zhang. 2021. Enhancing robustness
924 verification for deep neural networks via symbolic propagation. *Formal Aspects of Computing* 33, 3 (2021), 407–435.
- 925 [70] Meng Zhu, Weidong Min, Qi Wang, Song Zou, and Xinhao Chen. 2021. PFLU and FPFLU: Two novel non-monotonic activation functions in
926 convolutional neural networks. *Neurocomputing* 429 (2021), 110–117.