

Adversarial Robustness of Deep Neural Networks: A Survey from a Formal Verification Perspective

Mark Huasong Meng, Guangdong Bai, Sin Gee Teo, Zhe Hou,
Yan Xiao, Yun Lin, Jin Song Dong

Abstract—Neural networks have been widely applied in security applications such as spam and phishing detection, intrusion prevention, and malware detection. This black-box method, however, often has uncertainty and poor explainability in applications. Furthermore, neural networks themselves are often vulnerable to adversarial attacks. For those reasons, there is a high demand for trustworthy and rigorous methods to verify the robustness of neural network models. Adversarial robustness, which concerns the reliability of a neural network when dealing with maliciously manipulated inputs, is one of the hottest topics in cybersecurity and machine learning. In this work, we survey existing literature in adversarial robustness verification for neural networks and collect 39 diversified research works across machine learning, security, and software engineering domains. We systematically analyze their approaches, including how robustness is formulated, what verification techniques are used, and the strengths and limitations of each technique. We provide a taxonomy from a formal verification perspective for a comprehensive understanding of this topic. We classify the existing techniques based on property specification, problem reduction, and reasoning strategies. We also demonstrate representative techniques that have been applied in existing studies with a sample model. Finally, we discuss open questions for future research.

Index Terms—Robustness, verification, security, adversarial machine learning, neural networks, deep learning.

1 INTRODUCTION

DEEP learning is an *artificial intelligence* (AI) technique that is regarded as one of the top technological breakthroughs in computer science [1], [2]. It imitates the working principle of the human brain in processing data and forming knowledge patterns and produces promising results for various tasks such as image classification [3], [4], speech recognition [5], recommendation [6], and natural language understanding [7], [8]. Nowadays deep learning is increasingly applied in many fields where trustworthiness is critically needed, such as cybersecurity [9], autonomous driving [10], and the healthcare industry [11]. Particularly in cybersecurity, deep learning has been used to perform intrusion detection [12], malware detection [13], phishing detection [14], network traffic analysis [15], etc. Neural networks, or more specifically deep neural networks (DNNs), act as the key technology for deep learning. Most neural networks are designed to be deployed in practice where an arbitrary input shall be accepted. This, however, brings a great uncertainty in determining the behavior of a neural network model even if it has been well-trained [16]. As a result, the trustworthiness of neural networks has turned

out to be a new challenging topic in the research community, especially for those to be deployed in safety and security-critical systems [17].

The genesis of neural network verification and validation can be traced back to decades ago [18], in which the verification is referred to as correctness and the validation is referred to as accuracy and efficiency. Adversarial examples for neural network models were discovered by Szegedy *et al.* [19] and later Goodfellow *et al.* [20] systematically introduced an approach to intrude the robustness property of a neural network to make it misclassify a tampered input to a wrong result. Since then, neural networks are shown vulnerable to adversarial attacks [21]; consequently, research on the security of neural networks experienced explosive growth. Many research projects study the attack and defense of AI using adversarial training [22], [23], [24], [25]. As a result, verifying the robustness of neural network models against adversarial inputs becomes a strong demand in machine learning research to earn trust from their users and investors.

Thanks to a complete knowledge system of program analysis that has been built since the last century, many techniques designed for traditional program verification, such as *constraint solving* [26], *reachability analysis* [27], and *abstract interpretation* [28], are adopted in reasoning neural network models. On the other hand, researchers also attempt to reduce the verification to an optimization problem and achieve uplifting progress by taking advantage of various mathematics theories [29], [30], [31]. The progress made in robustness verification techniques not only increases industry practitioners' confidence in the security of neural

- M. H. Meng is with the National University of Singapore, and also with the Institute for Infocomm Research, A*STAR, Singapore. E-mail: menghs@i2r.a-star.edu.sg
- G. Bai is with The University of Queensland, Australia.
- S. G. Teo is with the Institute for Infocomm Research, A*STAR, Singapore.
- Z. Hou is with the Griffith University, Australia.
- Y. Xiao, Y. Lin and J. S. Dong are with the National University of Singapore.
- G. Bai (Email: g.bai@uq.edu.au) and Y. Xiao (Email: dcsxan@nus.edu.sg) are corresponding authors.

network models but also stimulates the discussion of more diverse property specifications in deep learning, such as *fairness* [32], *monotonicity* [33], and *coverage* [34].

An early survey under this theme [35] performs a comparative study on various verification methods, which collects seven different approaches from the literature and discusses the advantages as well as the shortcomings of each approach. A recent survey [36] extensively studies 18 existing robustness verification approaches and reproduces them on a unified mathematical framework, therefore for the first time comprehensively sketching the benchmarking of various verification techniques. However, some latest progress since 2019 has not been covered. Another survey [37] proposes a taxonomy of robustness verification in accordance with existing robust training techniques and conducts a unified evaluation toolbox for over 20 approaches. It takes both deterministic and probabilistic verification into account and unveils the landscape of available techniques for robustness training and robustness verification. Compared with existing literature, this paper covers the state-of-the-art robustness verification techniques that have been published in recent years. Moreover, we also propose a more comprehensive taxonomy to categorize the existing approaches and analyze them from a formal verification perspective, thereby forming our insight into future research on this topic.

Contributions

In summary, our paper makes the following contributions:

- This survey discusses adversarial robustness verification for deep neural networks from a perspective of formal verification including three key components: a property formalization, a reduction framework, and available reasoning strategies. From this point of view, we show the verification methodology not only differs in reasoning strategies but also varies with the author’s understanding of neural network verification as a problem.
- We propose a taxonomy based on different components under the formal verification framework and accordingly perform a classification of the existing literature. We establish the connection between different types of approaches from the technical perspective and thereby unveil the landscape as well as the latest progress. We particularly analyze the strategy behind the implementation of diversified approaches and discuss their advantages and shortcomings.
- Considering many existing approaches are proposed for specific use cases or experimental settings, it is hard to directly analyze them in a comparative context. This survey fills the gap by assessing them on a more general neural network model with a case study supported by substantial visualization.
- We suggest the future directions of deep neural network verification, including, improvement in completeness of verification, adoption of various types of complex models, leveraging optimization algorithms, and a generic verification framework.

Scope of This Survey

Considering the research on deep neural network property verification begins with the topic of adversarial robustness,

which is indeed the most discussed property so far, we choose robustness verification as the theme to carry out our survey. In order to avoid information overload and present this survey in an appropriate space, we only focus on reviewing and analyzing research works that put *robustness verification* as the major contribution.

We collect the existing literature published in machine learning, security, and software engineering domains in recent years. We also resort to two sources to replenish the literature. First, we take reference from the existing survey papers on neural network verification [35], [36], [37], and thereby we have collected some representative approaches. so that the representative approaches are not missed. Second, we take them as the seeds and treat the counterparts that have been cited and/or benchmarked in their papers as the second source. We note that we have only collected those that have been published; therefore, preprints that have not undergone peer review are excluded.

We also notice that most related studies simplify the context of a deep neural network as a *multilayered perceptron* (MLP) for classification tasks that are equipped with three typical activations (ReLU, sigmoid, and tanh), therefore we only discuss robustness verification of MLPs with three activations applicable in this survey. To avoid ambiguity, the term “neural networks” is referred to as deep neural networks with the assumption that it is in a multi-layer architecture rather than a single neuron perceptron. We also assume the robustness verification is performed in a *white-box* setting, which means the architecture and parameters of the examined model are transparent to the verifier.

As a result, we collect 39 research works that have proposed an algorithm or a tool to verify the robustness property, and we show them later in Table 1. For some research works that do not discuss robustness verification based on a predefined specification, but from the perspective of robust training or finding adversarial examples, we only take their methodologies that apply to specification-based verification into account.

2 BACKGROUND

This section introduces formal verification and one of its typical implementations called *model checking*, including its key elements and processes. After that, we propose our definition of property verification by fitting the setting of neural networks into the framework of formal verification.

2.1 Formal Verification

Formal verification is defined as a method to prove or disprove the correctness of a specific software or hardware system through formal methods of mathematics [38]. As traditional testing is commonly used for finding bugs through an actual execution or simulation, formal verification emphasizes proving the correctness in a modeling framework referring to the examined system with regard to a certain specification or property – and therefore formal verification requests an in-depth understanding of the system to be examined, and usually comes with a higher cost than traditional testing [39]. Model checking is a typical formal verification technique to systematically verify if a formally

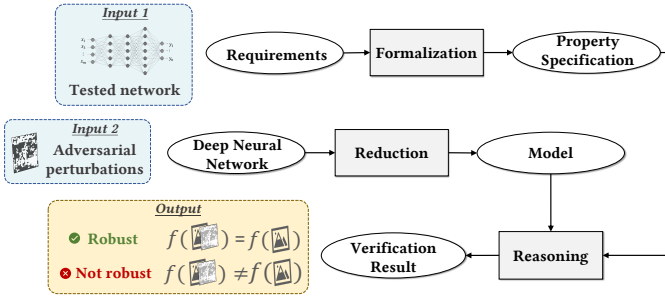


Fig. 1. An overview of adversarial robustness verification from a formal verification perspective

defined property holds in a finite state model [40]. It is designed to explore all possible system states exhaustively and, in the end, provides a qualitative result for specifying whether the target specification holds in the system. In Fig. 1 we present a classical schematic view of formal verification through model checking, which is composed of three key components: a *property formalization*, a *reduction framework*, and *specification reasoning*.

A precise and unambiguous property specification is necessary for formal verification. Property specification is a statement of property obtained through formalizing the system requirements. It is often presented as a *classical logic* formula to specify the relationship or restriction on different system components or illustrated in a *temporal logic* formula to enforce a rule over the system in terms of time. The reduction framework is a formal representation of the system to be examined, which keeps all the key features and functions to ensure a proper simulation of all possible states within the system. To sum up, the property formalization defines *what* the system is expected to do, and the system modeling after reduction prescribes *how* the system behaves. Property formalization and a proper reduction framework together constitute the *specification phase* of the verification. These two processes illustrate our understanding of system verification as a *problem*. Different understandings of a system may result in different reasoning approaches available for us to solve the problem.

With both system and property formalized during the specification phase, the verifier performs the reasoning task and outputs a result – we call this stage the *verification phase*. A positive output means that the property specification is satisfied in the examined model. However, a negative output sometimes could not be straightforwardly interpreted as a violation of that property specification, and therefore an extra analysis is required. Since the reasoning usually produces a counterexample together with a negative output, we need to determine if the counterexample generated through the simulation faithfully reflects the situation of the actual system – otherwise, it is concluded as a *false negative* and another round of reasoning is needed after *refining* the reduction framework. Such a strategy to recursively refine the system specification according to the counterexample generated from the previous round is formally proposed in [41] as the counterexample-guided abstraction refinement (CEGAR) approach.

2.2 Property Verification of Deep Neural Networks

After introducing the basic concept of neural networks and formal verification theory, we can define the problem of property verification of deep neural networks through the lens of formal verification. Here we treat a deep neural network model as a function that maps an input to an output, thereby we formalize the property specification as a logic formula involving both input and output. Given the property specification and available modeling framework on hand, we can perform the verification using diverse reasoning strategies.

Let us take the adversarial robustness as an example. Suppose we have a neural network $f : X \rightarrow Y$ that takes an input x and then outputs y as the result of classification. With a limit of the maximum input perturbation Δ_x given, we assert the neural network is not robust through the verification if we observe an arbitrary (x, y) from the training set such that $f(x + \Delta_x) \neq y$; otherwise the robustness is proved to be held in the given neural network.

Through the verification above, we can know if a property holds in the target model from the perspective of *the worst case*. However, we gain no knowledge of the percentage of compromised inputs (that are successfully manipulated by the attacker) within the entire input space, not even to mention the analysis of possible approaches to improve. For those reasons, *quantitative reasoning* is required to perform a more complex and in-depth property verification for neural network models. Counting the desirable outputs and offering a probabilistic result is a typical step of quantitative verification that makes it differ from qualitative verification. For those properties with demographic characteristics, such as fairness, quantitative reasoning could achieve better semantic expressiveness than any qualitative approach. Since we focus on robustness verification, we assume the property specification is qualitatively defined in this survey.

3 PROPERTY FORMALIZATION

The properties of a neural network could be captured from the semantic context of its specification. Most of the properties are *input-output (IO) properties* that request a specific mapping relation between the input and output. *Robustness* is one of the earliest IO properties that has been studied, which requests the model’s output to be stable upon minor modifications have been made to the input’s value [42], [43]. In this section, we present our definition of adversarial perturbations and provide a formalization of the robustness property.

3.1 An Introduction to Perturbation

The perturbation of input can be categorized into two types, namely *global perturbation* and *regional perturbation*, as illustrated in Fig. 2 (middle). As its name implies, the global perturbation is a vector of the same size as input samples, and the adversarial input is obtained by overlaying the perturbation on a benign input. The source of global perturbation can be from normal distortion in actual deployments such as signal noise or weather conditions, or the output of attack methods, such as *projected gradient descent* (PGD) [23] or *fast gradient sign method* (FGSM) [20]. Regional perturbation is usually a fixed pattern to be added to the original

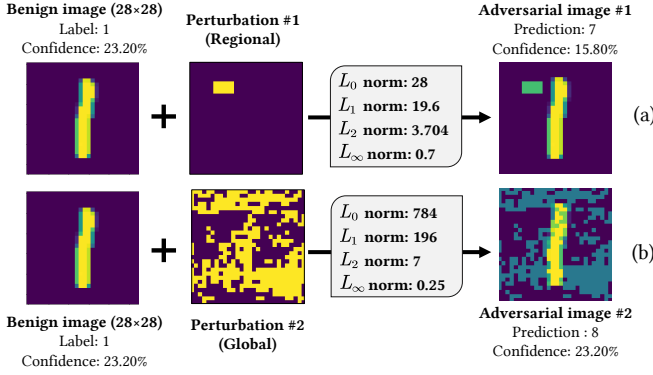


Fig. 2. Samples of regional (a) and global (b) adversarial perturbations and measurement of magnitude, illustrated with attacks to the prediction of the MNIST dataset

input, which is also known as local perturbation or “Trojan trigger” in an attack scenario. Due to the uncertainty of the occurrence of the attack, verification of security against Trojan attacks will be difficult without knowing the pattern, location, and size of triggers. Moreover, proving robustness against a specific Trojan pattern does not generalize the property against another Trojan pattern even when both patterns are on a similar scale. For those reasons, regional perturbation is mainly studied in attack scenarios. We focus on global perturbation in the remainder of this paper.

3.2 Perturbation Measurement

There are different ways to measure the magnitude (scope) of input perturbations. In mathematics, a *norm* is a function to map a vector’s distance to the origin into a non-negative value. Here we treat input perturbations as vectors, and we adopt the norm function to define the scope of perturbation, and correspondingly calculate its magnitude.

Although both global and regional perturbations are created with the same objective, which is to maximize the chance of misclassification meanwhile ensuring they are visually indistinguishable, the measuring methods for different types of perturbations may differ. L_0 norm calculates the number of input features that have been affected, i.e., the size of perturbation, which is useful to describe the size of a regional perturbation. L_1 norm (Manhattan distance) and L_2 norm (Euclidean distance) measure the distance between the benign and adversarial inputs, which are often used to describe the accumulated influence of the perturbation. L_∞ norm of a perturbation records the greatest magnitude among all elements of it as a vector. L_∞ norm is widely used in robustness verification because it echoes the ϵ value of the FGSM [20], which regulates the scope of gradient as the worst-case perturbation that leads to misclassification. Fig. 2 demonstrates a sample perturbation in each type given a benign input sample from the MNIST dataset [44] and illustrates the value of different norm functions in defining the magnitude of input perturbation.

We highlight that the adversarial examples are not the only source of input perturbations, which may also include benign corruption (e.g., the influence of weather) and reasonable distortion (e.g., signal noise, blur) [43]. However, those cases are specially introduced in object recognition applications and do not differ from the adversarial example

in the context of formal analysis. We assume adversarial examples as the default violation of robustness in this work.

3.3 Robustness against Adversarial Perturbations

Robustness describes the *fault-tolerant capability* of a computer system to deal with erroneous inputs during execution. In machine learning, robustness is usually described as the capacity of a model to consistently produce desirable output given some perturbations (or distortions) have been made on the input. As robustness is closely related to the reliability and performance of a neural network model, it is widely discussed in diverse branches of deep learning applications [43].

A mutant input with perturbation that can alter the model’s output is called an *adversarial example* – therefore the evaluation of the robustness of a deep neural network model is to find the existence of such an adversarial example. In analysis, robustness only makes sense if the scope of perturbation is clearly defined and the benign input is properly specified. The verification may not be able to truthfully reflect the real robustness of the given neural network model if the scope of perturbation is defined negligibly (i.e., too small). A too strong perturbation makes the adversarial example sufficiently differs from the original input, in which case it is reasonable for the model to classify the input as a different output. As for the benign input, it could be an image used for classification or a vector of values within a continuous domain. We then formalize the robustness property of neural network models as follows:

Definition 1 (Robustness against adversarial perturbations). *Given a neural network model that maps a benign input x_a to the output y_a , and moreover produces an output y_b given an adversarial input x_b . Let $\|\cdot\|_p$ be the function to calculate the L_p -norm distance. The robustness property Φ against any adversarial perturbation within the scope of Δ_x is defined as follows:*

$$\Phi(x_a, y_a, x_b, y_b, \Delta_x) \stackrel{\text{def}}{=} (\|x_a - x_b\|_p \leq \Delta_x) \rightarrow (y_a = y_b)$$

Other than the above definition that applies to the classification of a discrete domain, the robustness property could be further defined with *tolerant misclassification* or *tolerant misprediction* [33], which is more practical in the case that the output falls in a continuous domain (e.g., regression tasks) or an aggregation of labels (e.g., top-K classification). As this paper focuses on the classification scenario, we assume the verification does not tolerate errors in the output.

4 CORRECTNESS ASSESSMENT

From the view of program analysis, the objective of a static checker is to find if a program prevents something that it claims to avoid. The same idea also applies to the verification of neural network models. Intuitively, a positive verification outcome implies the neural network model could always prevent violations of specific property if it claims to do so and, on the contrary, a negative outcome claims the neural network model fails to preserve the robustness property. However, inconsistency between the claim and the fact sometimes happens when the verification is not guaranteed to be sound and complete. This section discusses *soundness* and *completeness* for a more comprehensive and precise correctness assessment [45].

		Hypothesis (Truth)	
		Desirables Prediction: 8 (ship) <i>Correct</i>	Violations Prediction: 0 (airplane) <i>Wrong</i>
Decision (Test)	Accepted i.e., the model is robust	Accepted Desirables (TRUE POSITIVE)	Missed Violations (FALSE NEGATIVE)
	Rejected i.e., the model is not robust	False Alarms (FALSE POSITIVE)	Caught Violations (TRUE NEGATIVE)

Fig. 3. The cases of verification outcomes classified in the statistical hypothesis testing approach, illustrated with the classification of the CIFAR-10 dataset

4.1 Soundness and Completeness of the Verification

A proof method is sound if every statement it can prove is indeed true. It is complete if every true statement can be proved. Given a neural network verifier with something undesired X that we wish to prevent. The verifier is sound if it never accepts X taking place in the neural network with certain inputs. It is complete if it never rejects any input, provided X does not happen in the neural network model.

Statistical hypothesis testing provides another good way to understand soundness and completeness, where the *hypothesis* represents the actual existence of an undesirable event X defined in the specification of a neural network model, and the *decision* indicates the actual outcome of the verifier [46]. A *table of error* specified for the verification is shown in Fig.3 with positive and negative hypothesis columns being replaced by desirables and violations. With such a framework being set up, we may understand the relationship between soundness and completeness as follows:

- A **sound** verification prevents *missed violations* (false negatives) but may tolerate *false alarms* (false positives) in case the event X does not actually take place even if the verifier outputs a rejection. Simply put, the verification may report *unsafe* regarding the (possible) existence of violation X , even though the system is indeed safe.
- A **complete** verification prevents false alarms (false positives). However, it may not be able to account for missed violations (false negatives) — which means there may exist some violations not found by the verification. The verifier may report the system is safe due to its failure to find the violation X .
- A **sound and complete** verification implies a *perfect prediction*, that only allows *accepted desirables* (true positives) and *caught violations* (true negatives).

4.2 Trade-off in Practice

Verifying deep neural network models is shown as an *NP-hard* problem in [26] and therefore it is challenging to implement verification that is simultaneously sound, complete, and able to terminate in a reasonable time.

In practice, one of those three requirements is usually put aside to guarantee the other two characters. Some approaches (e.g., [26], [47]) put aside termination to ensure both soundness and completeness are preserved, which is widely adopted in the early stage of neural network verification due to the size of neural networks is comparably small, and the scale of verification is limited. However, scalability




Fig. 4. The sample neural network (ReLU activation) with weight and bias parameters shown on edges and nodes respectively

becomes a critical factor in the verification of deeper and more complicated neural networks. Considering that soundness is more essential in program analysis (to prevent what the system claims to prevent), recent verification solutions (e.g., [48], [49]) commonly choose to sacrifice completeness to secure soundness and scalability.

Unlike the complete verification that comes with an exponential time complexity for the worst case, incomplete verifiers usually come with much better scalability but risk the verification from precision loss due to over-approximation. Since the precision loss is accumulated layer by layer, in the worst case, an incomplete verifier may fail to certify the robustness even though it terminates quickly.

4.3 Soundness of Floating-point Arithmetic

The soundness of floating-point arithmetic is another factor to determine the correctness of the verification. Floating-point is a commonly used data type in computing that offers flexible demands in precision. It does not naturally satisfy the axioms of real arithmetic, such as associativity and distributivity [50]. For that reason, the floating-point soundness is a concern in verifying neural network models on a larger scale. Unsound floating-point arithmetic can make the verification suffer from false negatives. At the moment of writing this survey, only a few papers [48], [50], [51] take floating-point arithmetic into consideration. Since we focus on studying robustness verification methodology from a formal verification perspective, floating-point arithmetic is not included in our classification criteria.

5 MODEL REDUCTION

This section discusses how to reduce the property verification of a neural network model to another problem with finite states and an available solution. Such a reduction process is also known as *modeling* in model checking. The verification reduction reflects our understanding of the verification as a problem and meanwhile determines what consequent approach or algorithm could be used to solve it. As one of the three key processes defined in Section 2.1, reduction stands for a formal interpretation of the neural network model to be examined.

We categorize the existing reduction methods into two groups, namely *optimization problem* and *reachability problem*. In case the verification is treated as an optimization problem, we further classify the reduction of existing approaches

thereby we can reduce the robustness verification to a constraint satisfaction problem, or more specifically a Boolean satisfiability problem (colloquially known as SAT problem).

Alternatively, we can also set the negation of constraints over the output nodes as the *objective function* and keep the remaining constraints unchanged, thus the verification process is transformed into solving an optimization problem. In the case of all activation are *piecewise-linear functions* (e.g., ReLU), this reduction approach is also known as linear programming (LP) encoding.

As we are checking whether a neural network is robust, the goal of reasoning over this encoding is to prove the existence of such a solution that satisfies all the constraints, rather than finding the objective value. To this end, we symbolize every node with a unique variable and then connect those variables layer-by-layer according to the computational relation defined in the network model. This reduction approach (in SAT form) can be formalized as below.

Definition 2 (SAT/LP encoding). *Given a neural network model f with n nodes. To verify a property specification of it, f can be formally reduced to a model \mathcal{M} in form of a constraint satisfaction problem, such as:*

$$\mathcal{M}_{(SAT/LP)} = \langle V, D, C \rangle$$

where $V = \{V_1, \dots, V_n\}$ is a set of variables corresponding to the nodes within f ; $D = \{D_1, \dots, D_n\}$ is a set specifying the value domain (if any) of variables in V ; and $C = \{C_1, \dots, C_m\}$ is the set of constraints, including the constraint reflecting the violation of property specification.

Fig.5 depicts the encoding of the sample network that violates the robustness specification, i.e., $x_{11} \leq x_{12}$. Because ReLU is a *piecewise-linear* function that produces different outputs from two input domains, we notice that each ReLU activation function introduces a disjunctive logical ($\circ r$) operation, which bisects the verification problem. As a result, given n ReLU nodes in a neural network model, its SAT/LP encoding creates a disjunction of 2^n conjunctive normal form (CNF) formulas (16 in the sample model) in total.

As we capture all nodes from the examined model and truthfully encode both affine and activation functions in constraints, we consider this reduction approach preserves both soundness and completeness for the next step of reasoning. Nevertheless, the nature of propositional logic is only suitable to encode *linear algebraic relations* without exponential. It makes this reduction approach only applicable to neural networks that use piecewise-linear activation, such as ReLU or its variants; otherwise, the output of activation must be relaxed by another linear or piecewise-linear function in advance.

5.2 Mixed-Integer Linear Programming (MILP) Encoding

One significant progress made on SAT/LP encoding is achieved by adding an integer variable for each ReLU activation and thereby encodes the examined model to a *mixed-integer linear programming* (MILP) problem. MILP is a variant of LP that allows part of variables specified as integers. Some recent papers [55], [57], [59], [67], [82] take advantage of the integer setting of MILP, propose a subtle encoding

$$\begin{aligned} V &= \{x_1, x_2, \dots, x_{12}\}, D = \{-1 \leq x_1 \leq 1, -1 \leq x_2 \leq 1\} \\ C &= \{x_3 = x_1 + x_2, x_4 = x_1 - x_2 - 1, \\ &((x_3 \leq 0) \wedge (x_5 = 0)) \vee ((x_3 > 0) \wedge (x_5 = x_3)), \\ &((x_4 \leq 0) \wedge (x_6 = 0)) \vee ((x_4 > 0) \wedge (x_6 = x_4)), \\ &x_7 = -x_5 + x_6 - 1, x_8 = x_5 - x_6 - 1, \\ &((x_7 \leq 0) \wedge (x_9 = 0)) \vee ((x_7 > 0) \wedge (x_9 = x_7)), \\ &((x_8 \leq 0) \wedge (x_{10} = 0)) \vee ((x_8 > 0) \wedge (x_{10} = x_8)), \\ &x_{11} = x_9 - x_{10} + 1, x_{12} = x_{10} - 0.5, x_{11} \leq x_{12} \} \end{aligned}$$

Fig. 5. Boolean formula obtained through SAT encoding of the sample network given the robustness specification being violated

of the ReLU function into a set of four in-equations, and thus eliminate the excessive disjunctive formulas in the constraints.

We show how a canonical integer setting could be adopted in ReLU activation in Definition 3 and present our formalization of MILP encoding in Definition 4. We then demonstrate the corresponding encoding of the sample network in Fig.6.

Definition 3 (Big-M encoding of ReLU). *Consider given an integer variable $d \in \{0, 1\}$ and another sufficiently great constant M (which is guaranteed to be greater than the upper bound of ReLU input), a ReLU function $y = \max(0, x)$ could be encoded by a CNF of four logical formulas, as shown below:*

$$\begin{aligned} y = \max(0, x) &\Leftrightarrow ((y \geq x) \wedge (y \leq x + M \cdot d) \\ &\wedge (y \geq 0) \wedge (y \leq M \cdot (1 - d))) \end{aligned}$$

where the case $d = 0$ represents the ReLU is in active mode ($x > 0$ and $y = x$), and conversely, $d = 1$ indicates the ReLU is inactive ($x < 0$ and $y = 0$).

Definition 4 (MILP encoding). *Given a neural network model f with n nodes. To verify a property of it, f can be formally modeled as a MILP problem presented as follows:*

$$\mathcal{M}_{(MILP)} = \langle V, V_{int}, D, C, Obj \rangle$$

where $V = \{V_1, \dots, V_n\}$ is a set of variables corresponding to the nodes within f ; V_{int} is the set of integers that correspondingly defined to constraint ReLU activation outputs; $D = \{D_1, \dots, D_n\}$ is a set specifying the value domain (if any) of variables in V ; $C = \{C_1, \dots, C_m\}$ is the set of constraints; and Obj is the objective (minimum) inequation looking for the counterexample in case of violation (i.e., $Obj \leq 0$) occurs.

It is worth noting that some existing approaches (e.g., [30], [77], [80], etc) initiate their analysis from an LP or MILP perspective and later transform the neural network encoding into a semidefinite programming (SDP) problem. We separate this encoding approach as an independent category and detail it in Section 5.3.

The MILP encoding is proposed as a successor of SAT/LP encoding. It retains both soundness and completeness at the reduction stage. Moreover, it addresses our concern about the ‘‘too many branches’’ issue in SAT/LP encoding, especially when handling larger or deeper neural networks, and eventually brings potential improvement regarding the scalability in neural network verification.

5.3 Quadratically Constrained Quadratic Program (QCQP) / Semidefinite Programming (SDP) Encoding

A quadratically constrained quadratic program (QCQP) is an optimization problem in which both the objective function and the constraints are quadratic functions. In fact, both LP and MILP encoding can be treated as variants of QCQP where the quadratic objective function and quadratic constraints are absent. Solving a QCQP is NP-hard, too. However, the problem can be solved by semidefinite programming (SDP) if it is convex. For that reason, existing studies categorized in this type are also known as SDP-based approaches [36], [60], [77]. They mainly transform the ReLU activation to convex quadratic equations, thereby the verification can be reduced to a convex QCQP after including the linear constraints (i.e., affine functions in the neural network). The quadratic encoding of ReLU can be defined as follows.

Definition 5 (QCQP encoding of ReLU). *Given a ReLU function $y = \max(0, x)$, we can encode it to a set of quadratic constraints, as shown below:*

$$y = \max(0, x) \Leftrightarrow ((y \geq x) \wedge (y \geq 0) \wedge (y(y - x) = 0))$$

Compared with LP and MILP encodings, the QCQP encoding enables us to analyze quadratic relations of symbolic (i.e., input and output of nodes) within a neural network model. By applying the convex encoding for ReLU activation, the QCQP can be solved with soundness and completeness guarantee using off-the-shelf interior point methods. The major concern with this type of encoding comes from scalability. Interior point methods are shown to be computationally expensive and could take up to an $O(n^6)$ complexity for a neural network model containing n hidden units [77].

Existing papers that adopt QCQP encoding commonly mitigate the scalability issue in two directions. They apply various mathematical optimization to improve solving efficiency and/or take advantage of the linear approximation to further relax the original QCQP. Particularly, recent work exploits linear approximation to extend the verification with a broader range of activation functions. Through the linear approximation (to be detailed in Section 6.4), both tanh and sigmoid functions can be approximately formulated as a set of quadratic equations.

5.4 Layerwise Propagation

The layerwise propagation faithfully and precisely simulates the computation of a neural network model to solve the verification as a reachability problem. It is intuitive and has been adopted in some early approaches [53], [58].

Given the property specification, we first identify the domain of concerned input (i.e., all the possible adversarial inputs). Then we simulate the calculation based on those inputs until reaching the output layer. In the end, we compare the propagated outputs with the domain of expected output to determine if the property preserves in the target model. We visualize the layerwise propagation process of verifying the sample model in Fig. 7-(a).

The layerwise propagation at the reduction stage guarantees both soundness and completeness as it in fact does

$$\begin{aligned} \text{minimize } & x_9 - 2 \times x_{10} + 1.5 \\ \text{s.t. } & x_1 + x_2 \leq x_5 \leq x_1 + x_2 + M \cdot d_3, \\ & x_5 \leq M \cdot (1 - d_3), \\ & x_1 - x_2 - 1 \leq x_6 \leq x_1 - x_2 - 1 + M \cdot d_4, \\ & x_6 \leq M \cdot (1 - d_4), \\ & x_5 + x_6 - 1 \leq x_9 \leq x_5 + x_6 - 1 + M \cdot d_7, \\ & x_9 \leq M \cdot (1 - d_7), \\ & x_5 - x_6 - 1 \leq x_{10} \leq x_5 - x_6 - 1 + M \cdot d_8, \\ & x_{10} \leq M \cdot (1 - d_8), \\ & -1 \leq x_1 \leq 1, \quad -1 \leq x_2 \leq 1, \\ & d_3, d_4, d_7, d_8 \in \{0, 1\}, \quad x_5, x_6, x_9, x_{10} \geq 0, \\ & M > x_i, \quad (1 \leq i \leq 12) \end{aligned}$$

Fig. 6. MILP encoding of the sample network after equations simplifying, with the objective expected to be positive

not explicitly alter the neural network as the body to be verified. The overall completeness and soundness are fully determined by the following reasoning strategies, which we will discuss in Section 6. Layerwise propagation usually works well on deep models, however, suffers while dealing with thick ones (i.e., lots of neurons per layer). Thus, the reasoning based on the layerwise propagation reduction without any optimization or trick would be computationally expensive, therefore leaving concern of scalability.

5.5 Discretization

The concern of scalability of layerwise propagation exists since the early stage of neural network verification. One possible approach to address such a concern is to mitigate the computational burden by *discretization*. Discretization is a concept in applied mathematics that describes the process of converting a continuous function into several discrete pieces, which aims to achieve a finite exploration within an infinite search space. The discretization has been adopted in a verification approach named DLV [56], which proposes to discretize the perturbation from the infinite input distribution. It propagates those discrete input vectors layer by layer and determines whether the robustness is preserved or not by observing the bounds of output. Fig. 7-(b) demonstrates the verification of the sample model through a discretization.

Although the architecture of the examined model remains unchanged, the domain of output completely depends on the discretized input domain. Unless the scale of input discretization has been set large enough to sufficiently depicts all the key points in the output domain, such as the *convex hull* of a hidden layer, there always exists a risk of precision loss and consequently leads to missed violation (false negative). Overall, discretization greatly simplifies the analysis, especially for a deep neural network that takes comparably lower dimension inputs. However, as a reduction process, it does not come with a guarantee of completeness. The guarantee of soundness is challenging, too, as it is mutually restricted by the granularity of discretization and the convexity of the computation at each layer. However, many neural networks, particularly those that use sigmoid or tanh activation, are not convex. For these reasons, this reduction approach may suffer from the state-space explosion issue and therefore does not scale to larger neural network models [28], [51].

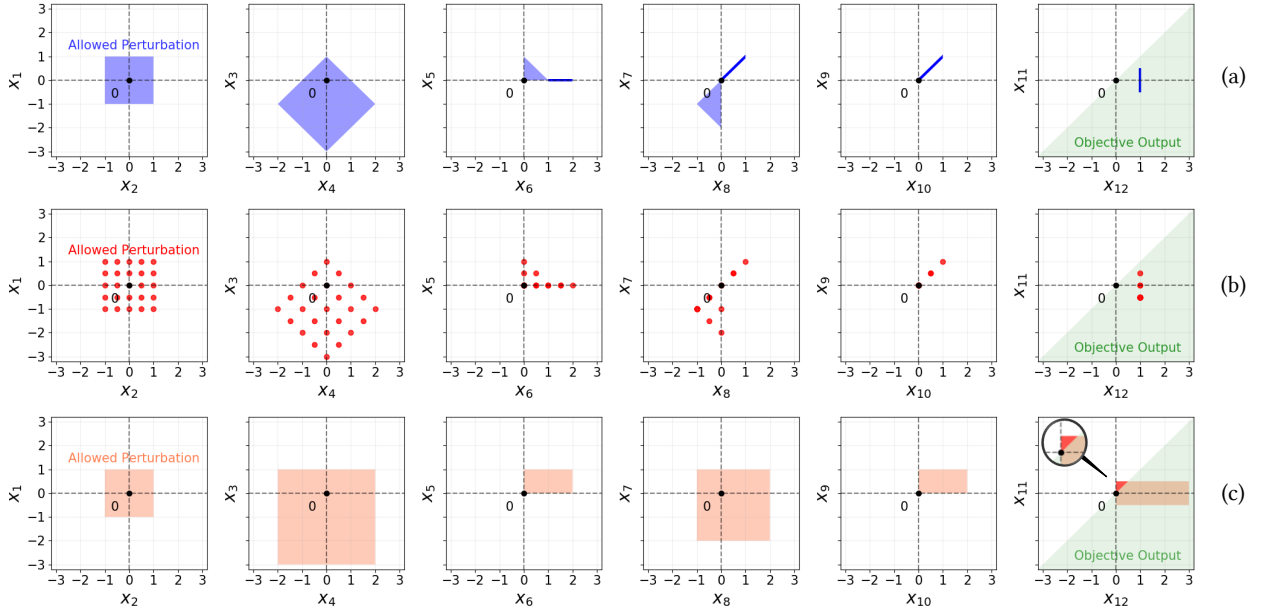


Fig. 7. Visualization of different reduction approaches in verifying robustness property of the sample model (abstraction interpretation with box domain and reasoned by interval arithmetic in sub-figure (a); discretization in sub-figure (b) and forward reachability analysis in sub-figure (c), with red color area indicating false positive produced)

5.6 Abstract Interpretation

Abstract interpretation is a verification approach that provides a formalism of approximation and abstraction in a mathematical setting, which has been adopted in [28], [49], [50] and their extensional works. An abstract interpretation consists of an *abstract domain*, a pair of *abstraction and concretization functions*, and a sound *abstract semantic function* in form of an iterative fixed-point equation.

Abstract interpretation aims to simplify the verification problem at the reduction stage. Although the verification is still treated as a reachability problem, it is solved by reasoning an abstract system rather than the original neural network model. The mainstream abstraction interpretation approaches symbolize each node of the neural network, apply abstraction function over its computation, and ultimately determine whether the property is preserved through the output values after the concretization. Besides that, [49] explores another direction of the abstract interpretation that simplifies the architecture of the examined model, rather than approximating the output of neurons inside.

In [49], each hidden layer has been simplified to contain at most four fully-connected neurons, literally one neuron with the maximum positive weight and tends to bring increment to the output written as $a^{[+, \uparrow]}$, one neuron with the minimize negative weight and tends to bring increment to the output $a^{[-, \uparrow]}$, one neuron with the maximum positive weight and tends to bring decrement to the output $a^{[+, \downarrow]}$, and lastly one neuron with the minimum negative weight and tends to bring decrement to the output $a^{[-, \downarrow]}$. Thus, both upper and lower bounds of the output can be eventually obtained by analyzing the abstract model. This approach gains a scalability advantage in verifying neural networks that are not very deep but built with a great number of hidden units per layer. On the other hand, simplifying neural network architecture through clustering the neurons' output does not

work well on models that adopt convergence functions as activation, such as sigmoid and tanh (because they always produce output within a certain range).

At the time of writing, DEEPPOLY [48] and its extensional approaches are the most comprehensive, precise, and practical solutions to perform verification by abstract interpretation. In the scheme of DEEPPOLY, each node is transformed into an abstract element. An abstract element is a tuple constituted by two *polyhedral constraints* and two *associate constants*. The polyhedral constraints of a node are made up of variables corresponding to those nodes connected from the previous layer, and the associate constants define its approximated upper and lower bounds after the concretization. Take the sample network as an instance, the robustness verification is reduced to determine if, in its concretized abstraction, the upper bound of x_{12} is less than the lower bound of x_{11} , which means $x_{11} > x_{12}$ is always true. We take [48] as an example and define the abstract interpretation reduction as follows.

Definition 6 (Abstract interpretation). *Given a neural network model f with n nodes. By defining a polyhedral abstract domain and approximating all the nodes accordingly, the verification of f can be achieved by concretizing its abstraction f_a , written as:*

$$\mathcal{M}_{(Abst.)} = f_a : \{\forall i \in [n] . a_i = \langle a_i^{\leq}, a_i^{\geq}, u_i, l_i \rangle\}$$

where a_i^{\leq} and a_i^{\geq} are polyhedral constraints of the abstract element corresponding to the i -th node, and u_i and l_i are associated constants that reflect the concretization of the i -th node's abstraction.

In Fig. 7-(c), we demonstrate the verification of the sample model by applying a box domain (each node is abstracted to a pair of lower and upper bounds) in the abstract interpretation. As we can observe from the output layer in the figure, the verification fails to determine whether the property holds due to a too coarse approxi-

mation (the tiny red area of the output domain *violates* the property specification). That shows the given property specification could not be verified in the sample model through this reduction approach, although it holds in the reality. Throughout the history of applying abstract interpretation in neural network verification, the abstract domain and corresponding abstraction functions have experienced evolution towards a more precise approximation of the neural network’s output, aiming to enhance the effectiveness of the verification through this reduction. Fortunately, the same property can be successfully verified by later works such as DEEPOLY [48], which we will discuss in the next section and visualize in Fig. 13.

Overall, abstract interpretation reduction enables analyzing the examined model’s behaviors in a less precise but more efficient way. Compared with the discretization that performs analysis merely based on discrete values of the input, abstract interpretation has an advantage in soundness guarantee as it ensures that all possible values of a given node are within the bounds obtained from over-approximation.

6 SPECIFICATION REASONING

Specification reasoning discusses in detail which tool or algorithm to be adopted in performing the verification. The choice of reasoning strategy closely relates to the reduction approach applied to the examined model and determines the overall verification performance. In this section, we classify available reasoning strategies into five classes and present the result in columns 15-19 of Table 1. With various reduction approaches defined in advance, we establish the connection between each reasoning method and its applicable reduction setting. Here we remark that one verification approach may adopt multiple specification reasoning strategies. Next, we detail each reasoning strategy and discuss its advantages and shortcomings.

6.1 Constraint Reasoning

Constraint reasoning is an analysis approach implemented by an off-the-shelf solver, and therefore it is also known as *solver reasoning*. This reasoning option is designed to take a set of constraints as input but does not restrict any reduction setting applied in advance. Reasoning through a satisfiability modulo theories (SMT) solver is the most intuitive way to perform verification of a neural network in SAT encoding. An SMT solver takes the SAT encoding of the examined model as input and heuristically searches for the best strategy to find a solution that satisfies all constraints.

Using an SMT solver to verify a large and deep neural network containing hundreds or even thousands of neurons is not an efficient approach. With the worst case of *exponential time complexity*, reasoning by an SMT solver usually cannot guarantee to terminate within a reasonable time. The high *cyclomatic complexity* constitutes another challenge. As each ReLU node generates a disjunction due to its piecewise-linear nature, reasoning a deep neural network with n ReLU neurons in hidden layers could produce 2^n sub-problems. Due to the lack of a polynomial-time algorithm for solving a CNF formula, constraint reasoning faces a

critical challenge of scalability. Dedicated tools using SMT solvers to solve CNF encoding are shown only able to handle small models containing 10-20 neurons in hidden layers [83].

One mitigation is replacing the SMT solver with an LP solver, which offers efficiency improvement in solving although sharing the same worst case of time complexity. For example, RELUPLEX [26] adopts an LP solver with an in-house implementation of the *simplex* algorithm in finding the feasibility of constraints. The simplex algorithm is known to be *nondeterministic polynomial (NP)* in time complexity but a remarkably effective way to solve an LP problem by a series of *update* and *pivot* operations. In RELUPLEX, a ReLU function is treated as a special equation. The value of ReLU output will be checked in each update operation. To avoid the infinite pivoting on variables within one ReLU function (either its input or output), RELUPLEX also adopts a branch-and-conquer strategy to split the original problem into two sub-problems, with each sub-problem corresponding to the *active* or *inactive* mode of that ReLU node. The experiment of [26] demonstrates a successful verification of 10 properties on the ACAS Xu¹ system. Owing to splitting heuristics, RELUPLEX could successfully verify neural network models with up to 300 ReLU nodes. It for the first time marks robustness verification *practical* in actual applications of deep neural networks. The latest progress of SMT/LP solver-based verification is MARABOU [66], which brings a significant improvement over its former release RELUPLEX by introducing a symbolic tightening algorithm and integrating an improved simplex algorithm.

Using a solver over the MILP encoding is another mitigation to address the scalability concern because it resolves the “branch explosion” issue in the SAT/LP encoding. *Gurobi* is the most used solver to reason the MILP encoding (see Table 2). Through the survey, we find [55], [57] might be the genesis of encoding a neural network into a MILP problem and reason by a solver. Later endeavors such as [59], [67], [82] also make use of this set of encoding and reasoning strategies to apply different solvers and attempt to support more types of neural network models. According to the benchmarking performed in [59], reasoning property specification as a MILP problem brings significant improvement in time performance. It allows us to analyze larger models containing up to 6,000 neurons.

Reasoning through a solver comes with an advantage that natively guarantees soundness and completeness, based on the premise that the solver is sound and complete. However, it faces scalability issues to a certain degree. Considering the desirable sound and complete verification can only be achieved if both modeling and reasoning approaches are sound and complete – at this moment, a solver-based verification, especially the MILP solver-based verification, is still the best choice to verify a deep neural network with a guarantee of completeness.

6.2 Mathematical Optimization

With the LP, MILP and SDP encoding selected, the verification of a neural network could also be achieved through

1. ACAS stands for Airborne Collision Avoidance System, where ACAS Xu is an experimental variant designed for the Remotely Piloted Aircraft Systems (RPAS).

TABLE 2

List of research works that disclose solver usage for the reasoning

Authorship, Year & Reference	Solver(s)
Pulina & Tacchella, 2010 (NEVER) [52]	HySAT
Scheibler <i>et al.</i> , 2015 [53]	iSAT
Cheng <i>et al.</i> , 2017 (MAXRESILIENCE) [55]	CPLEX
Ehler, 2017 (PLANET) [47]	miniSat, GLPK
Huang <i>et al.</i> , 2017 (DLV) [56]	Z3
Katz <i>et al.</i> , 2017 (RELUPLEX) [26]	GLPK
Lomuscio & Maganti, 2017 (NSVERIFY) [57]	Gurobi
Dutta <i>et al.</i> , 2018 (SHERLOCK) [59]	Gurobi
Wang <i>et al.</i> , 2018 (NEURIFY) [61]	Gurobi
Weng <i>et al.</i> , 2018 (FAST-LIN) [62]	Gurobi
Bunel <i>et al.</i> , 2018 (BAB, etc.) [65]	Gurobi
Katz <i>et al.</i> , 2019 (MARABOU) [66]	GLPK
Tjeng <i>et al.</i> , 2019 (MIPVERIFY) [67]	Cbc, CPLEX, Gurobi
Singh <i>et al.</i> , 2019 (REFINEZONO) [68]	Gurobi
Singh <i>et al.</i> , 2019 (KPOLY) [69]	Gurobi
Botoeva <i>et al.</i> , 2020 (VENUS) [76]	Gurobi
Bak <i>et al.</i> , 2020 (NNENUM) [72]	GLPK
Henriksen & Lomuscio, 2020 (VERINET) [75]	Gurobi

* Approaches that claim their usage of solvers but do not disclose the names are not listed.

various mathematical optimizations. Instead of generic programming (e.g., Python or C) with a specific constraint solver, approaches in this category are usually implemented based on off-the-shelf mathematical programming and numerical analysis platforms such as MATLAB.

One typical technique in this type is *duality*, which theoretically enables us to certify the robustness property by proving if the dual problem is solvable, although its original form is not practically easy to be solved. Wong and Kolter [29] proposed the first work that falls in this class. It illustrates the distortion over benign input as an adversarial polytope and applies a MILP encoding for the examined model. After that, it encodes the property violation into the *dual formulation* of the LP problem that corresponds to the original neural network model, and then applies *convex outer approximation* (also known as *relaxation*) in solving the dual problem. According to the duality theory, the original problem (also known as *primal*) can be guaranteed with a lower bound of the solution if its dual problem is found feasible to be solved. Moreover, when the number of variables blasts with the growth of the size of neural networks, the dual formulation becomes comparably easier to be solved than its primal. As a result, it provides a novel direction to make the intractable verification problem feasible to be solved through mathematical optimization.

Duality of LP/MILP only supports neural networks model with ReLU activation functions. Additionally, in case more non-piecewise-linear activation functions are supported or the examined neural network layers grow larger in size, the verification could suffer from false positives when the property fails to be certified – which owes to the gap between the dual and primal problems.

There is another approach proposed by Raghunathan *et al.* [30] that also adopts a similar idea of solving the dual of the original problem. In that work, the encoding of neural networks is presented as an SDP that is composed of *quadratic constraints*. Besides that, it presents the dual problem in the form of *Lagrange dual function* and uses

gradient descent to solve the optimization problem. Since all mainstream activation functions are differentiable, this approach supports a wide range of activation functions, including sigmoid and ReLU. Nevertheless, that work does not show good scalability, too. It only performs well on a fully-connected MLP with up to two layers. Recent work that applies a similar reasoning strategy to obtain the bound of the SDP problem includes [60], [77], [79], [80]. Another work by Dvijotham *et al.* [31] takes the limitation of previous research into account and proposes a general dual approach to support a broader class of neural networks. Compared with [29], [30], it claims to obtain non-trivial robustness bounds without the need for bound optimization during adversarial training. However, the quality of the robustness bound is traded off for the enhancement in scalability.

Reasoning by mathematical optimization is essentially introduced as an optimization technique in robust training. Unlike verifying a deep neural network model with a predefined property specification, this reasoning approach is more effective in finding the maximum robustness of a neural network model and helping in further neural network optimization. Reasoning by mathematical optimization faces challenges from both completeness and scalability. [79] improves verification scalability, however, only for neural networks with up to five layers. The disadvantage in scalability makes mathematical optimization less solely applied afterward. A recent approach [80] integrates mathematical optimization with linear approximation (to be detailed in Section 6.4) to enhance the overall efficiency, which presents a promising direction for future research.

6.3 Interval Arithmetic

Interval arithmetic, also known as interval analysis or bound propagation, is essentially developed as a classical method to measure the bounds of errors in mathematical computation [84]. It can also be used as an ideal method to evaluate the bounds of neural network output. To the extent of robustness specification reasoning, the preservation of robustness can be soundly certified once we establish the bounds of its output, which are propagated from the input perturbations, are within the acceptable range (i.e., the examined model does not misbehave).

Interval arithmetic has been adopted in early works [47], [55] as scalability mitigation of SAT/LP solver-based verification. By symbolizing all nodes in the examined neural network model, each node can be represented as a composition of variables at its previous layer. Given predefined bounds of the input, the output of the neural network can be evaluated through a layer-by-layer analysis. Considering both affine and ReLU functions are made up of simple arithmetic operations, we can follow the basic rules of interval operation to generate bounds for them [85].

Let $f : \mathbb{R}^J \rightarrow \mathbb{R}$ be the affine function of a neuron within hidden layers that executes the assignment $x_i \leftarrow b + \sum_{j \in [J]} (w_j \cdot a_j)$, where x_i has J neurons in the previous layer connected to it, and a_j stands for the output of the j -th neuron on the previous layer. Each a_j has a pair of the lower and upper bounds written as a_j^{\leq} and a_j^{\geq} . The interval of x_i , represented by $[x_i^{\leq}, x_i^{\geq}]$, can be obtained as follows:

$$x_i^{\leq} = b + \sum_{j \in [J]} \left(\max(0, w_j) \cdot a_j^{\leq} + \min(0, w_j) \cdot a_j^{\geq} \right)$$

$$x_i^{\geq} = b + \sum_{j \in [J]} \left(\max(0, w_j) \cdot a_j^{\geq} + \min(0, w_j) \cdot a_j^{\leq} \right)$$

Similarly, let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the ReLU activation function that executes the assignment $a_i \leftarrow \max(0, x_i)$, where x_i is the output of an affine function ahead and the input of the activation function. The bounds of a_i can be obtained as follows:

$$a_i^{\leq} = \max(0, x_i^{\leq}), \quad a_i^{\geq} = \max(0, x_i^{\geq})$$

We remark that interval arithmetic does not take symbolic dependency into account so it tends to generate a larger interval for a deep neural network when the number of layers increases. This phenomenon is also reflected in Fig. 7-(c) that the interval obtained at the output layer spans over the boundary of the objective (shown as the red color region), thus it fails to verify the robustness preservation. For that reason, even though interval arithmetic is simple and fast in computation, it can only provide a coarse evaluation of the examined neural network's output.

To mitigate the issue caused by too coarse evaluation, an iterative refinement (to be detailed in Section 6.5) sometimes is necessary to maintain the minimum bounds of neuron output in practice. For example, we can keep the symbolization along with layerwise propagation and then simplify the symbolic composition through backtracking to find a more precise interval of each neuron. Counterexample-based refinement is another useful strategy to refine the interval.

As a reasoning strategy, the usage of interval arithmetic is usually combined with linear approximation in recent literature as scalability becomes the primary issue in verifying larger and deeper neural networks. On the other hand, integrating interval arithmetic with a proper approximation method can also enable verification for those neural networks that use an activation without a piecewise-linear character like sigmoid and tanh functions.

6.4 Linear Approximation

In deep neural network verification, linear approximation (also known as linear relaxation) aims to provide an estimated range of non-linear activation output. It is firstly used as a complement to solver-based verification to extend its support with a broader range of activation functions [52], [83]. Linear approximation is one of the primary reasoning techniques for abstract interpretation modeling. By applying linear approximation with layerwise symbolic analysis, it is expected to get rid of reliance on SMT/MILP solvers and achieve a verification with provable soundness and scalability.

Next, we choose three typical activation functions that are used in most feedforward neural networks and present the existing linear approximation algorithms proposed in the literature. Some linear approximation algorithms also correspond to the abstraction function (e.g., zonotope [50] and polyhedron [48]) from the abstract interpretation perspective.




Fig. 8. ReLU function and its convex approximation of the upper bound (a) and both the upper and lower bounds (b), the shaded regions represent the range of value approximation.




Fig. 9. ReLU function and its convex approximation in DEEPPOLY for case (3) where $\lambda = 0$ (a) and $\lambda = 1$ (b), the shaded regions represent the range of value approximation.

ReLU Function

Linear approximation of the ReLU function is firstly proposed in [47] as a supplementary to the native LP modeling, which replaces the piecewise feature of the ReLU function by offering a pair of bounds, and thereby forms a constant mapping relation between the input and output of ReLU regardless of its activation status (i.e., active or inactive). Fig. 8-(a) illustrates an early linear approximation of ReLU with an upper bound. Another similar approximation that appears in [61], [62], [64] estimates both the upper and lower bounds of ReLU output in case the activation mode is unable to be determined (i.e., the input range spans over zero), as shown in Fig. 8-(b).

The state-of-the-art linear approximation of the ReLU function is proposed in DEEPPOLY [48] as the abstraction function over a polyhedral domain. Recall that in Definition 6 we propose a sound approximation of a node's value by two polyhedral constraints (i.e., a^{\leq} , a^{\geq}) and two associate constants (i.e., l , u). Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function that executes the assignment $x_i \leftarrow \max(0, x_j)$ for $j = i - 1$. Given $\langle a_j^{\leq}, a_j^{\geq}, l_j, u_j \rangle$, we can estimate the output in three cases:

- 1) If $l_j \geq 0$, $a_i^{\leq}(x) = a_i^{\geq}(x) = x_j$, $l'_i = l_j$, and $u'_i = u_j$;
- 2) If $u_j \leq 0$, $a_i^{\leq}(x) = a_i^{\geq}(x) = 0$, $l'_i = u'_i = 0$;
- 3) Otherwise $a_i^{\leq}(x) = \lambda \cdot x_j$, $a_i^{\geq}(x) = u_j \frac{(x_j - l_j)}{(u_j - l_j)}$, $l'_i = \lambda \cdot x_j$, $u'_i = u_j$ where $\lambda \in \{0, 1\}$ that minimizes the area of the resulting shape in the (x_i, x_j) -plane. i.e., if $|l_j| \geq |u_j|$, then $\lambda = 0$; otherwise $\lambda = 1$. Fig.9 illustrates two possible scenarios of the approximation in this case.

Compared with previous literature, the approximation proposed in DEEPPOLY produces the smallest approximation region (shown as shaded areas in Fig. 9), which stands for the highest precision among all relevant works.




Fig. 10. Interval abstraction proposed to evaluate the output of the sigmoid function, with input ranges over $[-2, 2]$ and $p = 0.5$.




Fig. 11. The convex approximation of the tanh function proposed in CROWN (a) and DEEPZ (b) when $l_j < 0 \leq u_j$, the shaded regions represent the range of over-approximation.

Sigmoid and Tanh Functions

Linear approximation of the sigmoid function is firstly proposed in [52] where the input and output of a sigmoid function (σ) are mapped with intervals. By doing that, a sigmoid function can be encoded into linear constraints containing input and output variables, and then be reasoned by an SMT/LP solver.

The output range of a sigmoid function is fixed as $(0, 1)$ and its derivation, $\sigma'(x) = \sigma(x)(1 - \sigma(x))$, has the maximum value of $1/4$ at which the input is zero. As shown in Fig.10, if we divide the input domain as a series of intervals in the length of p , we can always approximate the bounds of the sigmoid activation function with an arbitrary input x as $0 \leq \sigma(x+p) - \sigma(x) \leq p/4$. Thus we can initiate a sound abstract interpretation with a box abstract domain.

A similar method could also be applied to the tanh activation function since both of them have the maximum derivation value at zero points (the maximum derivation of $\tanh(x)$ is 1 where $x = 0$). This idea closely relates to another linear approximation by calculating the *Lipschitz continuous gradient* to assess output bounds, which becomes commonly adopted in later research such as [62], [70], [79].

One recent linear approximation technique for sigmoid and tanh functions is proposed in DEEPZ [50] as a part of zonotope abstract interpretation and later adopted in its succeeding works, such as DEEPPOLY [48]. Since both sigmoid and tanh functions present an S-shape on the input-output plane, they are always twice differentiable. Moreover, we observe their derivatives keep growing with the input value in the negative range, and conversely, declining with input's growth when the input becomes greater than zero. According to that character, we can estimate the growing slope of the output within a closed range determined by the lower and upper bounds of the input. This approximation is visualized in Fig. 12 and detailed as follows.

Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function that performs $x_i \leftarrow \sigma(x_j)$ or $x_i \leftarrow \tanh(x_j)$ for $j = i - 1$. Given $\langle a_j^{\leq}, a_j^{\geq}, l_j, u_j \rangle$,

we can approximate the new component i by setting two associate constants $l_i = g(l_j)$, $u_i = g(u_j)$, and two polyhedral constraints a_i^{\leq} and a_i^{\geq} according to two cases below:

- 1) If $u_j = l_j$, then $a_i^{\leq}(x) = a_i^{\geq}(x) = g(l_j)$;
- 2) Otherwise, we compute $a_i^{\leq}(x)$ and $a_i^{\geq}(x)$ separately.

We first set $\lambda = \frac{g(u_j) - g(l_j)}{u_j - l_j}$,

and $\lambda' = \min(g'(u_j), g'(l_j))$, then:

- (a) For the lower bound polyhedral constraint $a_i^{\leq}(x)$, we have:

$$a_i^{\leq}(x) = \begin{cases} g(l_j) + \lambda(x_j - l_j) & \text{if } l_j \geq 0 \\ g(l_j) + \lambda'(x_j - l_j) & \text{if } l_j < 0 \end{cases}$$

- (b) For the upper bound polyhedral constraint $a_i^{\geq}(x)$, we have:

$$a_i^{\geq}(x) = \begin{cases} g(u_j) + \lambda(x_j - u_j) & \text{if } u_j \leq 0 \\ g(u_j) + \lambda'(x_j - u_j) & \text{if } u_j > 0 \end{cases}$$

Another approximation is introduced in CROWN framework [64] and applied in later research such as [75]. It proposes a similar approximation of sigmoid and tanh functions with DEEPZ but offers more precise bounds when the input range spans over zero. Considering sigmoid and tanh functions always be *convex* given the input is less than zero, and be *concave* for the positive input, CROWN estimates the upper and lower bounds of sigmoid or tanh function at an arbitrary input point through a *tangent line* that originates from the input bounds, and thereby constructs a quadrilateral boundary.

Here we take the tanh function (written as g) as an example. Suppose the upper and lower bounds of input are given as $[l_j, u_j]$ such that $l_j < 0 < u_j$, we call the curve segment corresponding to the input bounds $g|_{[l_j, u_j]}$. There must exist a tangent line starting from either the lower bound or the upper bound of the tanh function touching the curve itself. As shown in Fig. 11(a), these two tangent lines constitute the upper and lower bounds of linear approximation. Since both tanh and sigmoid functions share the same shape of curve but only differentiate in output range, we omit the demonstration of the sigmoid function to save space.

Notwithstanding finding the tangent line for each bound could be computationally costly, it offers a sound estimation with higher precision. CROWN [64] proclaims the finding of tangent line can be achieved through a binary-search algorithm, which means a logarithmic time complexity on average and therefore theoretically does not escalate the overall difficulty of verification. Compared with the approximation proposed in CROWN [64], the zonotope abstraction of sigmoid and tanh functions proposed in DEEPZ [50] does not require finding tangent lines in calculating the output bounds. However, that approach produces a comparably coarse approximation (see Fig. 11), where the slope at either the lowest input or the highest input tends to be zero – in contrast CROWN [64] does not have such concern. Overall, the linear approximation strategy of DEEPZ [50] has the potential to achieve a better time efficiency in the practice but sacrifices precision to a certain degree.




Fig. 12. Visualization of the approximation of sigmoid/tanh activation that is proposed in DEEPZ based on its transformation in case (2) where $l_j \geq 0$ (a for sigmoid, d for tanh), $u_j < 0$ (b for sigmoid, e for tanh), and $l_j < 0 \leq u_j$ (c for sigmoid, f for tanh), the shaded regions represent the range of over-approximation.

6.5 Iterative Refinement

Iterative refinement has been proposed as a complementary technique to reinforce incomplete analysis such as over-approximation. Some existing approaches [48], [51], [72], [75] adopt it to achieve a *complete* verification despite diverse incomplete reasoning strategies being used, while most of the others take advantage of it to mitigate the incompleteness. It is firstly used in [52] to mitigate false positives in searching for the violation of a property specification. The iterative refinement strategy is inspired by the counterexample-guided abstraction refinement (CEGAR) in model checking. However, in the verification of neural networks, incomplete analysis usually does not offer a counterexample when the property specification is shown unsatisfied. In that case, the goal of iterative refinement is to adjust the “coarseness” of the approximation for certain nodes within the examined model and eventually validate if the unsatisfied result is caused by a counterexample or it is merely a false alarm.

Next, we categorize different techniques that have been used as iterative refinement by the existing approaches.

Constraint Replenishment

Constraint replenishment is widely-used when there is no parameter available to improve the precision, and therefore we have to bring more constraints for a less coarse analysis. A common way to achieve that is by adding a complete reasoning technique, for example referring the MILP encoding to the ReLU function to introduce additional constraints over its output, making use of these additional constraints in the feedforward analysis, and in the end obtaining a less coarse approximation of the neural network output [61], [68]. Recent literature such as [68], [72], [75], [76], [81] mainly perform linear approximation in the beginning, especially for the purpose of relaxing non-linear activation in a linear format, and then apply an off-the-shelf cross-platform solver library (e.g., Gurobi) to execute the verification.

Branch and Bound

Branch and bound (colloquially known as B&B, BaB) is an algorithm design paradigm for solving NP-hard optimization and interval analysis [86]. It is in nature a linear relaxation of non-linear computation for constraint reasoning. The branching takes advantage of an efficient *searching* strategy to prioritize the ReLU nodes for the analysis, followed by the bounding process that exploits the cutting-edge over-approximation techniques to find the bounds of the node.

Branch-and-bound is firstly adopted in [52] as a parameter justification technique to exclude unsatisfied verification outcomes caused by spurious counterexamples. Since the coarseness of the sigmoid approximation is determined by the size of an interval of p (see Fig. 10), higher precision can be achieved by heuristically assigning a smaller value to p . It repeats until we can determine if the property specification is satisfied, otherwise a counterexample that violates the property specification is given.

A similar strategy has also been applied in approximating the ReLU activation in [26], [47], [65], [72], [80] by bisecting its input domain. RELUPLEX [26] applies this strategy in their simplex algorithm to split each ReLU constraint into two constraints, i.e., the negative part that always leads to a zero output, and the positive part that remains unchanged through the ReLU computation. In [65], branch-and-bound has been adopted with multiple reasoning options to achieve completeness to different degrees. NNENUM [72] implements an enumerated searching algorithm in finding the sign of the ReLU input to refine bounds obtained from over-approximation and claims to achieve a complete verification in the end. Another recent work [80] also takes advantage of this approach to aid constraint reasoning, performing 50% better than the top candidates in recent verification competitions [87].

Multiple Neurons Analysis

Multiple neurons analysis is an emerging strategy that is proposed to enhance precision and mitigate incompleteness. It refines the spaces of each symbolic through a multiple variate analysis, in an intra-layer or cross-layer manner. It is firstly proposed in KPOLY [69] called *k-polyhedral convex hull approximation*, which reasons not only the constraints of feedforward propagation but also the linear relations of multiple neurons in the same layer. In addition to its previous work [48], KPOLY (by setting $k = 2$) calculates the bounds of sum ($x_i + x_j$) and difference ($x_i - x_j$) of two nodes at each layer. As shown in Fig. 13, by analyzing two nodes simultaneously, we can achieve higher precision than other approaches with only a single neuron analysis [48] along with the layer propagation. As a result, the size of bounding boxes at the output layer drops exponentially with the number of neurons to be analyzed at the same layer, implying less unnecessary precision loss produced due to the abstract interpretation.

Another recent paper [78] uses a similar strategy but with different reduction and reasoning strategies. It applies multivariate space refinement in both layerwise propagation




Fig. 13. Visualization of layerwise bounds refinement over the abstract interpretation of the sample neural network (finer bounds in orange color are obtained through interval arithmetic based on polyhedral abstract domain aided with backtracking in DEEPPOLY, and bounds in blue colors are further refined by applying multiple variate analysis that is adopted in KPOLY)

(FASTC2V) and LP solving (OPTC2V), and thereby achieves better completeness than previous approaches at the price of sacrifice in scalability.

Compared with other iterative refinement approaches, nowadays multiple neurons analysis still heavily relies on specific mathematical libraries or solving tools and therefore faces challenges to be adopted in generic programming or well-known machine learning frameworks (e.g., PyTorch, TensorFlow, etc.).

7 FUTURE DIRECTIONS

The upsurge of research in user-guide neural network verification started since the imperceptible input perturbation is proved to cause misclassification in several research works [19], [20] and more importantly, in real-life applications [88]. We find dozens of works published during the past decade that attempt to address this challenge and verify that a neural network model is robust against adversarial input perturbations. Next, we discuss a few potential research directions that are potential to bring progress in the future.

Endeavor for a better completeness

Our research community has been relentlessly working to strike a balance between completeness and scalability. To resolve this dilemma, one potential solution could be a heuristic combination of techniques used in both incomplete and complete verifiers. In one of the most recent works by Singh *et al.* [68], a methodology as a mixture of abstract interpretation and MILP solving is proposed, where two techniques jointly take effect during the verification. After the over-approximation is finished on each layer, a solver starts refining the boundary obtained from approximation with a timeout (because it is a complete verification method that does not guarantee a termination). That idea has preliminarily proved to gain an advantage from both complete and incomplete verifiers. However, due to the limitation of the state-of-the-art MILP approaches, only piecewise-linear activation (i.e., ReLU) is supported in that work, which means there is still much work to do in the future to implement a practical tool for the industry.

Adoption of various complex models

Most of the existing studies only perform the verification on MLP models, which is an elementary neural network architecture from nowadays view. As the max-pooling function could be treated as a variant of the ReLU function that outputs the maximum from multiple values, we are

pleased to see some recent studies (e.g., [28], [29], [48], [74]) have adopted their verification approaches to convolutional neural networks (CNNs). However, deep neural networks are not created only for image classification. Robustness is equally important for voice recognition and text prediction tasks, which are achieved by diverse complex models, such as recurrent neural networks (RNNs) and transformers. In that case, the challenges of soundness and completeness trade-off and support of diverse types of functions (not necessarily to be activation function) become critical and urgent. Recent research has made effort on extending existing reduction and reasoning techniques to robustness verification of RNNs [89], graph neural networks (GNNs) [90], L_∞ distance networks [91], variational auto-encoders (VAEs) [77], and transformers [92]. We foresee the existing approaches will be one day applied to more types of complex models.

Leveraging optimization algorithms

Existing LP/MILP encoding-based methods rely on mathematical optimization formulations and attempt to find the minimum value of an objective function with a large number of constraints. If the minimum value is positive, then the verification answers true. While the state-of-the-art is successful in exploring the optimization state-space using advanced branching and relaxation techniques [80], [81], there is still potential room for improvement. One possibility is to leverage recent advancements of nature-inspired optimization algorithms [93] at the reasoning stage. Such algorithms are highly efficient in finding optimal values of extremely complex objective functions at the cost that the results are often approximations rather than the global maximum/minimum. In the setting of the verification problem, this corresponds to trading completeness for scalability. However, with a customized algorithm dedicated to verification, it is possible to compute very close approximations of the minimum value for a large deep neural network in a significantly shorter time than existing methods. The above can be integrated into a method that combines incomplete and complete algorithms and takes advantage of both worlds à la [68]. It is also worth exploring the possibility of addressing different activation functions using such optimization algorithms. The outcome of these directions should be valuable despite being incomplete.

A generic verification framework

The advancement of verification techniques cannot benefit real-life AI applications unless it can be implemented in

generic programming frameworks and well supports mainstream machine learning ecosystems. We find that much literature in this field releases a proof-of-concept that can only work on particular models and mathematical programming platforms like MATLAB. Only around 30% (13 out of 39, see Table 1) of surveyed papers have implemented their approaches based on TensorFlow and/or PyTorch, which are the two most popular deep learning frameworks as of 2022 [94]. This phenomenon seriously limits the industrial adoption of robustness verification. We advocate that efforts need to be made for a more generic verification framework. Future research is suggested to release their techniques as public packages that are compatible with the mainstream deep learning ecosystems.

8 CONCLUSION

The verification of neural networks is the art of problem-solving. Its development engages persistent research efforts from verification, optimization, and machine learning communities. In this survey, we collect 39 existing approaches to verifying adversarial robustness and we propose our taxonomy for classifying those related works from a perspective of formal verification. We present our classification result regarding three key components of model checking including the formalization of property specification, reduction of the examined model, and available reasoning strategies. To gain an in-depth understanding of the rationality and difference between those diverse approaches, we create a sample multi-layered neural network model and provide substantial visualization in the paper. We then discuss both advantages and challenges that are yet to be addressed for each approach. In the end, we share our outlook on the future directions of deep neural network verification.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT press, 2016.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [4] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [5] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [6] S. Deng, L. Huang, G. Xu, X. Wu, and Z. Wu, "On deep learning for trust-aware recommendations in social networks," *IEEE transactions on neural networks and learning systems*, vol. 28, no. 5, pp. 1164–1177, 2016.
- [7] A. Bordes, S. Chopra, and J. Weston, "Question answering with subgraph embeddings," *arXiv preprint arXiv:1406.3676*, 2014.
- [8] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," *Advances in neural information processing systems*, vol. 27, 2014.
- [9] M. H. Vrejoiu, "Neural networks and deep learning in cyber security," *Romanian Cyber Security Journal*, vol. 1, no. 1, pp. 69–86, 2019.
- [10] Y. Tian, K. Pei, S. Jana, and B. Ray, "Deeptest: Automated testing of deep-neural-network-driven autonomous cars," in *Proceedings of the 40th international conference on software engineering*, 2018, pp. 303–314.
- [11] B. Alipanahi, A. Delong, M. T. Weirauch, and B. J. Frey, "Predicting the sequence specificities of dna-and rna-binding proteins by deep learning," *Nature biotechnology*, vol. 33, no. 8, pp. 831–838, 2015.
- [12] N. Shone, T. N. Ngoc, V. D. Phai, and Q. Shi, "A deep learning approach to network intrusion detection," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 1, pp. 41–50, 2018.
- [13] R. Nix and J. Zhang, "Classification of android apps and malware using deep neural networks," in *2017 International Joint Conference on Neural Networks (IJCNN)*, 2017, pp. 1871–1878.
- [14] K. Shima, D. Miyamoto, H. Abe, T. Ishihara, K. Okada, Y. Sekiya, H. Asai, and Y. Doi, "Classification of url bitstreams using bag of bytes," in *2018 21st Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, 2018, pp. 1–5.
- [15] Z. Okonkwo, E. Foo, Q. Li, and Z. Hou, "A cnn based encrypted network traffic classifier," in *Australasian Information Security Conference (AISC) 2022*, 2022.
- [16] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight uncertainty in neural networks," in *International conference on machine learning*. PMLR, 2015, pp. 1613–1622.
- [17] A. Shahrokni and R. Feldt, "A systematic review of software robustness," *Information and Software Technology*, vol. 55, no. 1, pp. 1–17, 2013.
- [18] L.-M. Fu, *Neural networks in computer intelligence*. Tata McGraw-Hill Education, 2003.
- [19] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [20] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *3rd International Conference on Learning Representations*, Y. Bengio and Y. LeCun, Eds., 2015.
- [21] M. Stewart, "Security vulnerabilities of neural networks," 2019, (Accessed 15-February-2022). [Online]. Available: <https://towardsdatascience.com/hacking-neural-networks-2b9f461ffe0b>
- [22] J. Jeong, S. Park, M. Kim, H.-C. Lee, D.-G. Kim, and J. Shin, "Smoothmix: Training confidence-calibrated smoothed classifiers for certified robustness," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [23] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *6th International Conference on Learning Representations*. OpenReview.net, 2018.
- [24] K. Leino, Z. Wang, and M. Fredrikson, "Globally-robust neural networks," in *International Conference on Machine Learning*. PMLR, 2021, pp. 6212–6222.
- [25] B. Zhang, T. Cai, Z. Lu, D. He, and L. Wang, "Towards certifying l-infinity robustness using neural networks with l-inf-dist neurons," in *International Conference on Machine Learning*. PMLR, 2021, pp. 12368–12379.
- [26] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient smt solver for verifying deep neural networks," in *International Conference on Computer Aided Verification*. Springer, 2017, pp. 97–117.
- [27] W. Ruan, X. Huang, and M. Kwiatkowska, "Reachability analysis of deep neural networks with provable guarantees," in *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, J. Lang, Ed. ijcai.org, 2018, pp. 2651–2659.
- [28] T. Gehr, M. Mirman, D. Drachler-Cohen, P. Tsankov, S. Chaudhuri, and M. Vechev, "Ai2: Safety and robustness certification of neural networks with abstract interpretation," in *2018 IEEE Symposium on Security and Privacy*. IEEE, 2018, pp. 3–18.
- [29] E. Wong and J. Z. Kolter, "Provable defenses against adversarial examples via the convex outer adversarial polytope," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. G. Dy and A. Krause, Eds., vol. 80. PMLR, 2018, pp. 5283–5292.
- [30] A. Raghunathan, J. Steinhardt, and P. Liang, "Certified defenses against adversarial examples," in *6th International Conference on Learning Representations*. OpenReview.net, 2018.
- [31] K. Dvijotham, R. Stanforth, S. Goyal, T. A. Mann, and P. Kohli, "A dual approach to scalable verification of deep networks," in *The Conference on Uncertainty in Artificial Intelligence*, vol. 1, 2018, p. 2.
- [32] A. Albarghouthi, L. D'Antoni, S. Drews, and A. V. Nori, "Fairsquare: probabilistic verification of program fairness," *Proceedings of the ACM on Programming Languages*, vol. 1, no. OOPSLA, pp. 1–30, 2017.

- [33] S. A. Seshia, A. Desai, T. Dreossi, D. J. Fremont, S. Ghosh, E. Kim, S. Shivakumar, M. Vazquez-Chanlatte, and X. Yue, "Formal specification for deep neural networks," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2018, pp. 20–34.
- [34] Y. Sun, M. Wu, W. Ruan, X. Huang, M. Kwiatkowska, and D. Kroening, "Concolic testing for deep neural networks," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 109–119.
- [35] W. Xiang, P. Musau, A. A. Wild, D. M. Lopez, N. Hamilton, X. Yang, J. Rosenfeld, and T. T. Johnson, "Verification for machine learning, autonomy, and neural networks survey," *arXiv preprint arXiv:1810.01989*, 2018.
- [36] C. Liu, T. Arnon, C. Lazarus, C. Strong, C. Barrett, M. J. Kochenderfer *et al.*, "Algorithms for verifying deep neural networks," *Foundations and Trends® in Optimization*, vol. 4, 2020.
- [37] L. Li, X. Qi, T. Xie, and B. Li, "Sok: Certified robustness for deep neural networks," *arXiv preprint arXiv:2009.04131*, 2020.
- [38] A. Sanghavi, "What is formal verification?" May 2010, (Accessed 15-August-2020). [Online]. Available: https://archive.eetasia.com/www.eetasia.com/ART_8800607430_480100_TA_7e47bbe1.htm
- [39] H. Wayne, "Why don't people use formal methods?" January 2019, (Accessed 15-January-2022). [Online]. Available: <https://www.hillelwayne.com/post/why-dont-people-use-formal-methods>
- [40] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press, 2008.
- [41] E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith, "Counterexample-guided abstraction refinement," in *International Conference on Computer Aided Verification*. Springer, 2000, pp. 154–169.
- [42] N. Carlini and D. Wagner, "Towards Evaluating the Robustness of Neural Networks," in *2017 IEEE Symposium on Security and Privacy*. San Jose, CA, USA: IEEE, May 2017, pp. 39–57.
- [43] D. Hendrycks and T. G. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," in *7th International Conference on Learning Representations*. OpenReview.net, 2019.
- [44] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [45] M. Hicks, "What is soundness (in static analysis)?" October 2017, (Accessed 20-January-2022). [Online]. Available: <http://www.pl-enthusiast.net/2017/10/23/what-is-soundness-in-static-analysis/>
- [46] B. Meyer, "Soundness and completeness: With precision," April 2019, (Accessed 19-January-2022). [Online]. Available: <https://cacm.acm.org/blogs/blog-cacm/236068-soundness-and-completeness-with-precision>
- [47] R. Ehlers, "Formal verification of piece-wise linear feed-forward neural networks," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2017, pp. 269–286.
- [48] G. Singh, T. Gehr, M. Püschel, and M. Vechev, "An abstract domain for certifying neural networks," *Proceedings of the ACM on Programming Languages*, vol. 3, no. POPL, pp. 1–30, 2019.
- [49] Y. Yisrael Elboher, J. Gottschlich, and G. Katz, "An abstraction-based framework for neural network verification," in *Computer Aided Verification*. Springer International Publishing, 2020, pp. 43–65.
- [50] G. Singh, T. Gehr, M. Mirman, M. Püschel, and M. Vechev, "Fast and effective robustness certification," in *Advances in Neural Information Processing Systems*, 2018, pp. 10 802–10 813.
- [51] S. Wang, K. Pei, J. Whitehouse, J. Yang, and S. Jana, "Formal security analysis of neural networks using symbolic intervals," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 1599–1614.
- [52] L. Pulina and A. Tacchella, "An abstraction-refinement approach to verification of artificial neural networks," in *International Conference on Computer Aided Verification*. Springer, 2010, pp. 243–257.
- [53] K. Scheibler, L. Winterer, R. Wimmer, and B. Becker, "Towards verification of artificial neural networks." in *MBMV*, 2015, pp. 30–40.
- [54] O. Bastani, Y. Ioannou, L. Lampropoulos, D. Vytiniotis, A. Nori, and A. Criminisi, "Measuring neural net robustness with constraints," in *Advances in neural information processing systems*, 2016, pp. 2613–2621.
- [55] C.-H. Cheng, G. Nührenberg, and H. Ruess, "Maximum resilience of artificial neural networks," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2017, pp. 251–268.
- [56] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu, "Safety verification of deep neural networks," in *International Conference on Computer Aided Verification*. Springer, 2017, pp. 3–29.
- [57] A. Lomuscio and L. Maganti, "An approach to reachability analysis for feed-forward ReLU neural networks," *ArXiv*, vol. abs/1706.07351, 2017.
- [58] W. Xiang, H.-D. Tran, and T. T. Johnson, "Reachable set computation and safety verification for neural networks with relu activations," *arXiv preprint arXiv:1712.08163*, 2017.
- [59] S. Dutta, S. Jha, S. Sankaranarayanan, and A. Tiwari, "Output range analysis for deep feedforward neural networks," in *NASA Formal Methods Symposium*. Springer, 2018, pp. 121–138.
- [60] A. Raghunathan, J. Steinhardt, and P. S. Liang, "Semidefinite relaxations for certifying robustness to adversarial examples," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [61] S. Wang, K. Pei, J. Whitehouse, J. Yang, and S. Jana, "Efficient formal safety analysis of neural networks," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 6367–6377.
- [62] L. Weng, H. Zhang, H. Chen, Z. Song, C.-J. Hsieh, L. Daniel, D. Boning, and I. Dhillon, "Towards fast computation of certified robustness for relu networks," in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 5276–5285.
- [63] W. Xiang, H.-D. Tran, and T. T. Johnson, "Output reachable set estimation and verification for multilayer neural networks," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 11, pp. 5777–5783, 2018, publisher: IEEE.
- [64] H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel, "Efficient neural network robustness certification with general activation functions," in *Advances in neural information processing systems*, 2018, pp. 4939–4948.
- [65] R. Bunel, I. Turkaslan, P. H. Torr, P. Kohli, and M. P. Kumar, "A unified view of piecewise linear neural network verification," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 4795–4804.
- [66] G. Katz, D. A. Huang, D. Ibeling, K. Julian, C. Lazarus, R. Lim, P. Shah, S. Thakoor, H. Wu, A. Zeljić *et al.*, "The marabou framework for verification and analysis of deep neural networks," in *International Conference on Computer Aided Verification*. Springer, 2019, pp. 443–452.
- [67] V. Tjeng, K. Xiao, and R. Tedrake, "Evaluating robustness of neural networks with mixed integer programming," in *7th International Conference on Learning Representations*. OpenReview.net, 2019.
- [68] G. Singh, T. Gehr, M. Püschel, and M. T. Vechev, "Boosting robustness certification of neural networks," in *7th International Conference on Learning Representations*, 2019.
- [69] G. Singh, R. Ganvir, M. Püschel, and M. Vechev, "Beyond the single neuron convex barrier for neural network certification," in *Advances in Neural Information Processing Systems*, 2019, pp. 15 098–15 109.
- [70] C. Huang, J. Fan, W. Li, X. Chen, and Q. Zhu, "Reachnn: Reachability analysis of neural-network controlled systems," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 18, no. 5s, pp. 1–22, 2019.
- [71] W. Xiang, H.-D. Tran, X. Yang, and T. T. Johnson, "Reachable set estimation for neural network control systems: A simulation-guided approach," *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [72] S. Bak, H.-D. Tran, K. Hobbs, and T. T. Johnson, "Improved geometric path enumeration for verifying relu neural networks," in *Proceedings of the 32nd International Conference on Computer Aided Verification*. Springer, 2020.
- [73] H.-D. Tran, D. Manzananas Lopez, P. Musau, X. Yang, L. V. Nguyen, W. Xiang, and T. T. Johnson, "Star-based reachability analysis of deep neural networks," in *International symposium on formal methods*. Springer, 2019, pp. 670–686.
- [74] H.-D. Tran, X. Yang, D. Manzananas Lopez, P. Musau, L. V. Nguyen, W. Xiang, S. Bak, and T. T. Johnson, "Nnv: the neural network verification tool for deep neural networks and learning-enabled cyber-physical systems," in *International Conference on Computer Aided Verification*. Springer, 2020, pp. 3–17.

- [75] P. Henriksen and A. Lomuscio, "Efficient neural network verification via adaptive refinement and adversarial search," in *ECAI 2020*. IOS Press, 2020, pp. 2513–2520.
- [76] E. Botoeva, P. Kouvaros, J. Kronqvist, A. Lomuscio, and R. Misener, "Efficient verification of relu-based neural networks via dependency analysis," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 3291–3299.
- [77] S. Dathathri, K. Dvijotham, A. Kurakin, A. Raghunathan, J. Uesato, R. R. Bunel, S. Shankar, J. Steinhardt, I. Goodfellow, P. S. Liang, and P. Kohli, "Enabling certification of verification-agnostic networks via memory-efficient semidefinite programming," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 5318–5331.
- [78] C. Tjandraatmadja, R. Anderson, J. Huchette, W. Ma, K. K. PATEL, and J. P. Vielma, "The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [79] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety verification and robustness analysis of neural networks via quadratic constraints and semidefinite programming," *IEEE Transactions on Automatic Control*, 2020.
- [80] B. Batten, P. Kouvaros, A. Lomuscio, and Y. Zheng, "Efficient neural network verification via layer-based semidefinite relaxations and linear cuts," in *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, Z.-H. Zhou, Ed. International Joint Conferences on Artificial Intelligence Organization, 8 2021, pp. 2184–2190, main Track.
- [81] P. Kouvaros and A. Lomuscio, "Towards scalable complete verification of relu neural networks via dependency-based branching," in *International Joint Conference on Artificial Intelligence*, 2021, pp. 2643–2650.
- [82] M. Fischetti and J. Jo, "Deep neural networks and mixed integer linear optimization," *Constraints*, vol. 23, no. 3, pp. 296–309, 2018, publisher: Springer.
- [83] L. Pulina and A. Tacchella, "Challenging smt solvers to verify neural networks," *Ai Communications*, vol. 25, no. 2, pp. 117–135, 2012.
- [84] H. Dawood, *Theories of interval arithmetic: mathematical foundations and applications*. LAP Lambert Academic Publishing, 2011.
- [85] G. Alefeld and G. Mayer, "Interval analysis: theory and applications," *Journal of computational and applied mathematics*, vol. 121, no. 1-2, pp. 421–464, 2000.
- [86] J. Clausen, "Branch and bound algorithms-principles and examples," *Department of Computer Science, University of Copenhagen*, pp. 1–30, 1999.
- [87] VNN-COMP, "Verification of neural networks competition (vnn-comp 2020)," 2020, (Accessed 15-January-2022). [Online]. Available: <https://sites.google.com/view/vnn20/vnncomp>
- [88] L. Alan and B. Ryan, "Tesla's autopilot found partly to blame for 2018 crash on the 405," <https://www.latimes.com/business/story/2019-09-04/tesla-autopilot-is-found-partly-to-blame-for-2018-freeway-crash>, 2019, (Accessed 1-February-2022).
- [89] X. Du, Y. Li, X. Xie, L. Ma, Y. Liu, and J. Zhao, "MARBLE: Model-based robustness analysis of stateful deep learning systems," in *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. ACM, Dec. 2020, pp. 423–435.
- [90] B. Wang, J. Jia, X. Cao, and N. Z. Gong, "Certified robustness of graph neural networks against adversarial structural perturbation," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 1645–1653.
- [91] B. Zhang, D. Jiang, D. He, and L. Wang, "Boosting the certified robustness of l-infinity distance nets," *arXiv preprint arXiv:2110.06850*, 2021.
- [92] Z. Shi, H. Zhang, K.-W. Chang, M. Huang, and C.-J. Hsieh, "Robustness verification for transformers," in *International Conference on Learning Representations*, 2019.
- [93] S. Mirjalili, J. S. Dong, and A. Lewis, Eds., *Nature-Inspired Optimizers - Theories, Literature Reviews and Applications*, ser. Studies in Computational Intelligence. Springer, 2020, vol. 811.
- [94] R. O'Connor, "Pytorch vs TensorFlow in 2022," <https://www.assemblyai.com/blog/pytorch-vs-tensorflow-in-2022/>, 2021, (Accessed 11-April-2022).



Mark Huasong Meng is a Research Engineer at Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR), Singapore. He received his bachelor's degree at Nanyang Technological University in 2014, and master's degree at the National University of Singapore in 2016. His research interests include trustworthy AI, software engineering, and system security.



Guangdong Bai received his bachelor's and master's degrees in computer science from Peking University, China, in 2008 and 2011, respectively, and his Ph.D. degree in computer science from the National University of Singapore in 2015. He is now a Senior Lecturer with the University of Queensland. His research interests include cyber security, protocol verification, and software engineering.



Sin Gee Teo is a Scientist at Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR), Singapore. He obtained the Ph.D. degree from Monash University, Australia in 2016. His research interests include applied cryptography, data privacy and security, malware and network anomaly classification, federated learning, and deep learning.



Zhe Hou is a Lecturer at Griffith University, Australia. He obtained his Ph.D. degree from the Australian National University in 2015. His research mainly focuses on automated reasoning, formal methods, autonomous system, machine learning, and blockchain.



Yan Xiao is a Research Fellow at the National University of Singapore. She received her Ph.D. degree from the City University of Hong Kong in 2019. Her research focuses on the trustworthiness of deep learning systems and AI applications in software engineering.



Yun Lin received his Ph.D. from Fudan University. He is currently a Research Assistant Professor at the National University of Singapore. His research mainly focuses on program analysis, traditional software testing and quality assurance analysis of artificial intelligence. He has published some top-tier conference/journal papers relevant to software analysis in ICSE, ISSTA, FSE, ASE, TIFS, TSE, AAAI, and USENIX Security.



Jin Song Dong received his Ph.D. degree from the University of Queensland, Australia. He is a professor with the School of Computing, National University of Singapore. His research interests include software engineering, program analysis, formal verification, and model checking.