Novelty Assessment Report

Paper: Sound Verification of Deployed Neural Networks

PDF URL: https://openreview.net/pdf?id=5IHuwRwMHK

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-07

Abstract

Verification methods aim at mathematically proving desirable properties of neural networks, such as robustness to adversarial perturbations. A verifier is sound if and only if it never claims that a neural network has the desired property when it does not. It was shown recently that none of the currently known verifiers that are claimed to be sound are guaranteed to be sound when considering the deployed version of the verified network. Due to this, all the known verifiers are vulnerable to certain backdoor attacks, where an adversarial network passes verification but, in reality, it exhibits adversarial behavior in specific deployment environments. So far, it has been suspected that sound verification is prohibitively expensive if we wish to verify all possible executions—including parallel and stochastic ones—in deployment. We are the first to propose an efficient error bounding technique that most known verifiers can apply to become practically sound. The technique enables both interval bound propagation and symbolic propagation methods to remain sound even if the deployment environment randomly selects a valid ordering and parenthesizing of the arithmetic operations to compute the network. We present a theoretical foundation for our approach and demonstrate empirically that our technique indeed discovers all known deployment-specific attacks, introducing only a limited performance overhead.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Sound Verification of Neural Networks under Floating-Point Arithmetic**

A total of **30 papers** were analyzed and organized into a taxonomy with **21 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Floating-Point Error Analysis and Modeling**
- **Sound Verification Approaches and Tools**
- **Robustness Analysis under Floating-Point Constraints**
- **Certified Proof Production and Trust**
- **Quantization and Mixed-Precision Optimization**
- **Theoretical Foundations and Universal Approximation**
- **Hardware Implementation and Floating-Point Arithmetic**

Complete Taxonomy Tree

- Sound Verification of Neural Networks under Floating-Point Arithmetic Survey Taxonomy
- Floating-Point Error Analysis and Modeling
    - Backward Error Analysis for Neural Networks (2 papers)
    - [1] Backward error analysis of artificial neural networks with applications to floating-point computations and adversarial attacks (Beuzeville, 2024) View paper
    - [2] Deterministic and probabilistic backward error analysis of neural networks in floating-point arithmetic (Buttari, 2024) View paper
    - Probabilistic Floating-Point Error Analysis (2 papers)
    - [5] Rigorous Roundoff Error Analysis of Probabilistic Floating-Point Computations (George M. Constantinides, 2021) View paper
    - [6] Rigorous Roundoff Error Analysis of Probabilistic Floating-Point\n Computations (Constantinides George, 2021) View paper
    - Floating-Point Accumulation Network Verification (1 papers)
    - [4] Automatic Verification of Floating-Point Accumulation Networks (Zhang, 2025) View paper
    - Floating-Point Approximation and Constraint Propagation (1 papers)
    - [29] Correct Approximation of IEEE 754 Floating-Point Arithmetic for Program Verification (Bagnara, 2019) View paper
- Sound Verification Approaches and Tools
    - Sound Deployment-Aware Verification ★ (2 papers)
    - [0] Sound Verification of Deployed Neural Networks (Anon et al., 2026) View paper
    - [11] No Soundness in the Real World: On the Challenges of the Verification of Deployed Neural Networks (Banhelyi, 2025) View paper
    - Interval-Based and Symbolic Propagation Verification (3 papers)
    - [16] Scalable Neural Network Geometric Robustness Validation via HÃ¶lder Optimisation (Y Zhang, 2025) View paper
    - [19] Sound Floating-Point Neural Network Verification with MILP (Shifu Yang, 2024) View paper
    - [23] Scaling up the static analysis of neural networks using affine forms (Soualah, 2024) View paper
    - SMT-Based Verification with Quantization (1 papers)
    - [27] QNNVerifier: A Tool for Verifying Neural Networks using SMT-Based Model Checking (Xidan Song, 2021) View paper
    - Input-Space Quantization Verification (1 papers)
    - [18] Verifying Low-dimensional Input Neural Networks via Input Quantization (Jia Kai, 2021) View paper

  - Software-Level Floating-Point Verification (1 papers)
  - [22] Floating-Point Neural Network Verification at the Software Level (Manino, 2025) View paper
  - Nondeterminism-Aware Optimistic Verification (1 papers)
  - [20] Nondeterminism-Aware Optimistic Verification for Floating-Point Neural Networks (Yao, 2025) View paper
- Robustness Analysis under Floating-Point Constraints
  - Randomized Smoothing Soundness (1 papers)
  - [3] Sound randomized smoothing in floating-point arithmetics (VorÃ¡Äek, 2022) View paper
  - Adversarial Robustness and Floating-Point Exploits (2 papers)
  - [8] Exploiting Verified Neural Networks via Floating Point Numerical Error (Jia Kai, 2021) View paper
  - [15] Fooling a complete neural network verifier (Banhelyi, 2021) View paper
  - Formal Robustness Verification Surveys (2 papers)
  - [9] Verification of Neural Networks (RÃ¶ssig, 2020) View paper
  - [30] Taxonomy and Techniques: Formal Verification of Adversarial Robustness in Deep Learning (W Aya, n.d.) View paper
- Certified Proof Production and Trust
  - Proof-Producing Neural Network Verifiers (3 papers)
  - [7] A Certified Proof Checker for Deep Neural Network Verification (Desmartin, 2025) View paper
  - [10] Neural network verification with proof production (Isac, 2022) View paper
  - [25] A Certified Proof Checker for Deep Neural Network Verification in Imandra (Desmartin, 2024) View paper
  - Theorem Proving for Floating-Point Verification (1 papers)
  - [26] Taming Floating-Point Rounding Errors with Proofs (Invited Talk) (Titolo, 2025) View paper
- Quantization and Mixed-Precision Optimization
  - Rigorous Mixed-Precision Tuning (1 papers)
  - [14] Rigorous floating-point mixed-precision tuning (Wei-Fan Chiang, 2017) View paper
  - Quantization Verification with Classification Guarantees (1 papers)
  - [24] Quantization with Guaranteed Floating-Point Neural Network Classifications (Anan Kabaha, 2025) View paper
- Theoretical Foundations and Universal Approximation
  - Floating-Point Universal Approximation Theory (1 papers)
  - [12] Floating-Point Neural Networks are Provably Robust Universal Approximators (Hwang, 2025) View paper
  - Comprehensive Verification Frameworks (2 papers)
  - [17] Trustworthy Machine Learning for High Assurance Systems (Stell, 2025) View paper
  - [28] Certified Deep Learning: Verification and Training (Mirman, 2022) View paper
- Hardware Implementation and Floating-Point Arithmetic
  - FPGA-Based Neural Network Implementation (1 papers)
  - [13] 32-Bit Fixed and Floating-Point Hardware Implementation for Enhanced Inverter Control: Leveraging FPGA in Recurrent Neural Network Applications (Chanakya Hingu, 2024) View paper
  - Sound Floating-Point Analysis Tools (1 papers)
  - [21] Sound Analysis of Floating-Point Programs (Verduzco, 2023) View paper

Narrative

Core task: sound verification of neural networks under floating-point arithmetic. The field addresses the gap between idealized real-number semantics and the finite-precision arithmetic used in deployed systems. The taxonomy organizes research into several main branches: Floating-Point Error Analysis and Modeling examines how rounding errors propagate through network layers, with works like Backward Error Analysis[1] and Probabilistic Backward Error[2] developing formal frameworks for quantifying these deviations. Sound Verification Approaches and Tools focuses on building verifiers that account for floating-point semantics, including efforts such as Sound MILP Verification[19] and Software Level Verification[22] that ensure correctness at the implementation level. Robustness Analysis under Floating-Point Constraints investigates how finite precision affects adversarial robustness guarantees, exemplified by Randomized Smoothing Floating[3]. Certified Proof Production and Trust emphasizes generating machine-checkable certificates, as seen in Certified Proof Checker[7] and Imandra Proof Checker[25]. Quantization and Mixed-Precision Optimization explores reduced-precision representations, while Theoretical Foundations and Universal Approximation and Hardware Implementation branches address foundational questions and platform-specific concerns.

A central tension runs through these branches: balancing soundness guarantees with practical scalability. Many studies tackle the challenge of taming rounding errors in verification workflows, with works like Rigorous Roundoff Error[5] and Taming Rounding Errors[26] proposing rigorous yet tractable error bounds. The original paper, Sound Verification Deployed[0], sits within the Sound Deployment-Aware Verification cluster, emphasizing verification that reflects real deployment conditions rather than idealized models. This positions it closely alongside No Soundness Real[11], which critiques verification approaches that ignore implementation-level discrepancies, and contrasts with more abstract robustness frameworks like Randomized Smoothing Floating[3]. The deployment-aware perspective highlights open questions about bridging the gap between verification at the algorithmic level and guarantees that hold on actual hardware, a theme echoed across certified proof production efforts and hardware-aware quantization studies.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. No Soundness in the Real World: On the Challenges of the Verification of Deployed Neural Networks

Authors: Banhelyi, Balazs, Attila Sz'asz, Jelasity, Mark, et al. (7 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

The ultimate goal of verification is to guarantee the safety of deployed neural networks. Here, we claim that all the state-of-the-art verifiers we are aware of fail to reach this goal. Our key insight is that theoretical soundness (bounding the full-precision output while computing with floating point) does not imply practical soundness (bounding the floating point output in a potentially stochastic environment). We prove this observation for the approaches that are currently used to achieve pr...

Relationship Analysis

Both papers belong to the Sound Deployment-Aware Verification category, addressing the challenge of ensuring verification soundness across all possible execution orderings and deployment environments in neural networks under floating-point arithmetic. The original paper proposes a practical solution by developing efficient error bounding techniques that enable existing verifiers (IBP and symbolic propagation) to achieve practical soundness by covering all possible expression trees with limited overhead. In contrast, the candidate paper focuses on demonstrating the fundamental problem itself, proving theoretically that current verifiers are not practically sound and constructing deployment-specific backdoor attacks to empirically expose these vulnerabilities, without proposing a comprehensive solution to achieve practical soundness.

Contributions Analysis

Overall novelty summary. The paper proposes an efficient error bounding technique enabling existing verifiers to achieve soundness under floating-point arithmetic in deployed neural networks. It resides in the 'Sound Deployment-Aware Verification' leaf, which contains only two papers total, indicating a relatively sparse research direction within the broader taxonomy of sound verification under floating-point constraints. This positioning suggests the work addresses a recognized but underexplored gap: ensuring verification guarantees hold across all possible execution orderings and environments in real deployment, not just idealized mathematical models.

The taxonomy reveals that neighboring leaves focus on interval-based and symbolic propagation methods, SMT-based verification with quantization, and software-level floating-point verification. These adjacent directions emphasize algorithmic soundness or specific verification paradigms, whereas the deployment-aware cluster explicitly targets the gap between verification-time assumptions and deployment-time realities. The scope note for this leaf highlights verification of 'all possible execution orderings and environments,' distinguishing it from theoretical soundness approaches that may not account for parallel or stochastic execution contexts encountered in practice.

Among thirty candidates examined, the contribution-level analysis shows mixed results. The efficient error bounding technique examined ten candidates with zero refutations, suggesting novelty in the specific method proposed. However, the theoretical foundation for deployment-sound verification examined ten candidates and found one refutable match, indicating some overlap with prior theoretical work in this limited search scope. The two sound verification algorithms with empirical validation examined ten candidates with no refutations, pointing to potential novelty in the algorithmic instantiation and experimental validation aspects.

Based on the top-thirty semantic matches examined, the work appears to occupy a sparsely populated research direction with some theoretical overlap but distinct methodological contributions. The analysis does not cover the full breadth of verification literature, and the single refutation among thirty candidates suggests the theoretical foundation builds on recognized prior work while the algorithmic and empirical components may offer more distinctive advances within the deployment-aware verification paradigm.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Efficient error bounding technique for sound verification

Description: The authors introduce a novel bounding technique that enables verifiers to remain sound even when deployment environments randomly select valid orderings and parenthesizations of arithmetic operations. This technique allows both interval bound propagation and symbolic propagation methods to cover all possible expression trees in deployment.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. General Cutting Planes for Bound-Propagation-Based Neural Network Verification

URL: View paper

Brief Assessment

General Cutting Planes[34] focuses on bound propagation methods for neural network verification with cutting plane constraints, not on error bounding techniques for handling floating-point arithmetic variations across deployment environments.

2. A sound abstraction method towards efficient neural networks verification

URL: View paper

Brief Assessment

Sound Abstraction Method[37] appears to focus on abstraction methods for neural network verification but provides insufficient context to assess overlap with the original paper's error bounding technique for handling arbitrary expression trees in floating-point arithmetic.

3. Robustness verification of neural networks using polynomial optimization

URL: View paper

Brief Assessment

Polynomial Optimization Robustness[33] focuses on robustness verification using polynomial optimization and semidefinite programming techniques, not on error bounding for floating-point arithmetic operations across different expression trees in deployment environments.

4. Residual-based error bound for physics-informed neural networks

URL: View paper

Brief Assessment

Residual Error Bound[40] focuses on error bounds for physics-informed neural networks solving differential equations, not on sound verification of neural networks against adversarial perturbations or deployment-specific expression trees.

5. Branch and Bound for Piecewise Linear Neural Network Verification

URL: View paper

Brief Assessment

Piecewise Linear Branch[32] focuses on branch-and-bound methods for verifying piecewise linear neural networks through mixed integer programming formulations, not on error bounding techniques for handling floating-point arithmetic variations across deployment environments.

6. Exponent Relaxation of Polynomial Zonotopes and Its Applications in Formal Neural Network Verification

URL: View paper

Brief Assessment

Polynomial Zonotopes Exponent[38] focuses on restructuring polynomial zonotopes to reduce complexity in neural network verification, not on error bounding techniques for handling deployment-specific arithmetic operation orderings and parenthesizations.

7. Probabilistic verification of neural networks using branch and bound

URL: View paper

Brief Assessment

Probabilistic Branch Bound[31] focuses on probabilistic verification of neural networks using branch and bound with probability distributions over inputs, not on error bounding techniques for handling floating-point arithmetic expression trees in deployment environments.

8. Quantitative verification of neural networks and its security applications

URL: View paper

Brief Assessment

Quantitative Security Verification[36] focuses on probabilistic and quantitative reasoning for neural network properties (counting inputs satisfying properties), not on error bounding techniques for handling floating-point arithmetic and expression trees in deployment environments.

9. Bound Tightening Using Rolling-Horizon Decomposition for Neural Network Verification

URL: View paper

Brief Assessment

Rolling Horizon Decomposition[35] focuses on mixed-integer programming decomposition for neural network verification, not on error bounding techniques for handling floating-point arithmetic non-associativity and expression tree variations in deployment environments.

10. Bridging neural ode and resnet: A formal error bound for safety verification

URL: View paper

Brief Assessment

Neural ODE ResNet[39] focuses on bounding approximation errors between Neural ODE and ResNet models for safety verification of continuous-time systems, not on handling deployment-specific floating-point arithmetic variations and expression trees that the original paper addresses.

Contribution 2: Theoretical foundation for deployment-sound verification

Description: The paper provides formal proofs establishing that their bounding method correctly over-approximates the range of ReLU networks across all possible expression trees. This includes propositions and corollaries demonstrating soundness for both IBP and symbolic propagation approaches.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Framework design of Network intrusion detection based on convolutional neural networks

URL: View paper

Brief Assessment

Network Intrusion Detection[51] focuses on convolutional neural networks for intrusion detection systems, not on formal verification of ReLU networks or deployment-sound verification methods. The candidate addresses a completely different application domain (cybersecurity) with different technical objectives than the original paper's verification theory.

2. An abstract domain for certifying neural networks

URL: View paper

Brief Assessment

Abstract Domain Certifying[57] focuses on abstract domain design for neural network certification using floating point polyhedra and intervals, but does not address deployment-sound verification across different expression trees or floating-point environments as defined in the original paper.

3. No Soundness in the Real World: On the Challenges of the Verification of Deployed Neural Networks

URL: View paper

Prior Art Analysis

No Soundness Real[11] demonstrates that the original paper's theoretical foundation addresses a problem that was already identified and formalized in prior work. The candidate paper proves that theoretically sound verifiers (those that bound full-precision outputs) are not necessarily practically sound (do not bound deployed network outputs), and introduces formal definitions of deployed models and deployed verification problems. This prior work establishes the theoretical framework for distinguishing between theoretical and practical soundness before the original paper's contribution.

Evidence

Evidence 1 - **Rationale**: The candidate paper establishes the theoretical distinction between theoretical and practical soundness, which is the foundational problem that the original paper claims to address. This shows prior theoretical work on the same verification soundness problem. - **Original**: we are the first to propose an efficient error bounding technique that most known verifiers can apply to become practically sound. the technique enables both interval bound propagation and symbolic propagation methods to remain sound even if the deployment environment randomly selects a valid orderi... - **Candidate**: we formally prove that a verifier that is theoretically sound (i.e., bounds the full-precision output correctly) is not necessarily practically sound, that is, it might not bound the actual output correctly in a given deployment environment. the problem has practical implications.

Evidence 2 - **Rationale**: The original paper cites sz´asz et al. (2025) as prior work that argued theoretical and deployed models are different mathematical objects. The candidate paper (No Soundness Real[11]) is this cited work, proving the theoretical foundation before the original paper. - **Original**: sz´asz et al. (2025) went a step further, arguing that the theoretical and deployed (practical) models are different mathematical objects, and verification must consider the details of the deployment environment. - **Candidate**: our key insight is that theoretical soundness (bounding the full-precision output while computing with floating point) does not imply practical soundness (bounding the floating point output in a potentially stochastic environment). we prove this observation for the approaches that are currently used...

Evidence 3 - **Rationale**: Both papers use identical mathematical formulations for the deployed verification problem, with the candidate paper presenting this formulation first, indicating prior theoretical work on formalizing deployment-sound verification. - **Original**: for an input point x∗ ∈ x, the deployed verification problem is ∀x ∈ dϵ,p,e(x∗) ∀z ∈ r(x; θ, e), z≥ 0, (2) where dp,ϵ,e(x∗) = {x ∈ x : ‖x - x∗‖p ≤ ϵ}. - **Candidate**: given an input x∗ ∈ x, the verification problem is now to prove that ∀x ∈ dϵ,p,e(x∗) ∀z ∈ r(x; θ, e), z≥ 0, (2) where dp,ϵ,e is defined below. the problem is formulated assuming a stochastic environment, but note that the deterministic environment is a special case of the stochastic one.

4. Branch and Bound for Piecewise Linear Neural Network Verification

URL: View paper

Brief Assessment

Piecewise Linear Branch[32] provides theoretical foundations for branch-and-bound verification of piecewise linear networks but does not address deployment-sound verification across different expression trees or floating-point environments as claimed in the original contribution.

5. Seev: Synthesis with efficient exact verification for relu neural barrier functions

URL: View paper

Brief Assessment

Seev Synthesis Verification[56] focuses on verification of ReLU neural barrier functions for control systems, not deployment-sound verification across different floating-point environments and expression trees as addressed in the original paper.

6. Global optimization of objective functions represented by ReLU networks

URL: View paper

Brief Assessment

Global ReLU Optimization[53] focuses on global optimization of objective functions represented by ReLU networks, not on deployment-sound verification or floating-point soundness across different expression trees.

7. Complexity of Injectivity and Verification of ReLU Neural Networks (Extended Abstract)

URL: View paper

Brief Assessment

The candidate paper (Injectivity Complexity ReLU[58]) focuses on complexity analysis of injectivity verification for ReLU networks. No full text context is provided to assess overlap with deployment-sound verification theory.

8. Generating and checking dnn verification proofs

URL: View paper

Brief Assessment

Generating Checking Proofs[52] focuses on proof generation and checking for DNN verification tools using MILP formulations and activation pattern trees. It does not address theoretical foundations for deployment-sound verification of ReLU networks across different floating-point expression trees, which is the core contribution of the original paper.

9. Improved geometric path enumeration for verifying relu neural networks

URL: View paper

Brief Assessment

Geometric Path Enumeration[55] focuses on improving the efficiency of exact path enumeration for ReLU networks through algorithmic optimizations, not on establishing theoretical foundations for deployment-sound verification considering floating-point arithmetic and expression trees.

10. OSIP: Tightened Bound Propagation for the Verification of ReLU Neural Networks

URL: View paper

Brief Assessment

Cannot assess refutation as OSIP Tightened Bound[54] provides no full text context for comparison with the original paper's theoretical proofs for deployment-sound verification.

Contribution 3: Two sound verification algorithms with empirical validation

Description: The authors implement two verification algorithms (FPSoundIBP and FPSoundSymbolic) that incorporate their bounding technique. They prove these algorithms are practically sound and demonstrate empirically that they correctly detect all known deployment-specific attacks with limited performance overhead.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Efficient neural network verification via adaptive refinement and adversarial search

URL: View paper

Brief Assessment

Adaptive Refinement Search[43] focuses on symbolic interval propagation with adaptive splitting for ReLU/sigmoid/tanh networks, not on deployment-specific floating-point soundness. The original paper addresses floating-point arithmetic across all possible expression trees in deployment environments, which is a fundamentally different problem than the candidate's focus on efficient verification through adaptive refinement strategies.

2. Dynamic Back-Substitution in Bound-Propagation-Based Neural Network Verification

URL: View paper

Brief Assessment

Dynamic Back Substitution[41] focuses on improving computational efficiency of bound propagation methods without addressing floating-point soundness issues that are central to the original paper's contribution. The candidate develops dynamic back-substitution for standard verification settings, while the original addresses deployment-specific soundness guarantees accounting for all possible expression trees in floating-point arithmetic.

3. Achieving Verified Robustness to Symbol Substitutions via Interval Bound Propagation

URL: View paper

Brief Assessment

Symbol Substitutions Robustness[47] focuses on verified robustness for NLP text classification against symbol substitutions using interval bound propagation, not general neural network verification across deployment environments with floating-point arithmetic considerations.

4. Eager Falsification for Accelerating Robustness Verification of Deep Neural Networks

URL: View paper

Brief Assessment

Eager Falsification Acceleration[49] focuses on accelerating verification through eager falsification and subproblem ordering, not on developing sound verification algorithms that handle floating-point deployment issues. The candidate integrates with existing verifiers (MIPVerify, Neurify, DeepZ, DeepPoly) rather than proposing new sound algorithms like FPSoundIBP and FPSoundSymbolic.

5. Verification of Neural Control Barrier Functions with Symbolic Derivative Bounds Propagation

URL: View paper

Brief Assessment

Control Barrier Functions[44] focuses on verifying neural control barrier functions for robot safety using symbolic derivative bounds propagation, not general neural network verification with interval bound propagation and symbolic propagation as in the original paper's deployment-specific verification context.

6. Optimized Symbolic Interval Propagation for Neural Network Verification

URL: View paper

Brief Assessment

Optimized Symbolic Interval[46] focuses on symbolic interval propagation optimization for neural network verification, not on sound verification algorithms addressing floating-point deployment issues or deployment-specific attacks as in the original paper.

7. Symbolic-numeric programming in scientific computing

URL: View paper

Brief Assessment

Symbolic Numeric Programming[45] focuses on symbolic-numeric programming systems for scientific computing, not neural network verification algorithms or interval bound propagation methods.

8. Formal security analysis of neural networks using symbolic intervals

URL: View paper

Brief Assessment

Symbolic Intervals Security[42] focuses on interval bound propagation and symbolic propagation for neural network verification but does not address deployment-specific soundness issues like floating-point expression trees that the original paper targets.

9. Robust Training of Neural Networks against Bias Field Perturbations

URL: View paper

Brief Assessment

Bias Field Perturbations[50] focuses on robust training methods using interval bound propagation and symbolic interval propagation for bias field perturbations in image classification, not on sound verification algorithms for deployment-specific attacks in neural networks.

10. Formal techniques for verification and testing of cyber-physical systems

URL: View paper

Brief Assessment

Cyber Physical Verification[48] focuses on formal verification techniques for cyber-physical systems and controller code analysis, not neural network verification algorithms using interval bound propagation and symbolic propagation for deployment-specific soundness.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Sound Verification of Deployed Neural Networks View paper
- [1] Backward error analysis of artificial neural networks with applications to floating-point computations and adversarial attacks View paper
- [2] Deterministic and probabilistic backward error analysis of neural networks in floating-point arithmetic View paper
- [3] Sound randomized smoothing in floating-point arithmetics View paper
- [4] Automatic Verification of Floating-Point Accumulation Networks View paper
- [5] Rigorous Roundoff Error Analysis of Probabilistic Floating-Point Computations View paper
- [6] Rigorous Roundoff Error Analysis of Probabilistic Floating-Point\n Computations View paper
- [7] A Certified Proof Checker for Deep Neural Network Verification View paper
- [8] Exploiting Verified Neural Networks via Floating Point Numerical Error View paper
- [9] Verification of Neural Networks View paper
- [10] Neural network verification with proof production View paper
- [11] No Soundness in the Real World: On the Challenges of the Verification of Deployed Neural Networks View paper
- [12] Floating-Point Neural Networks are Provably Robust Universal Approximators View paper
- [13] 32-Bit Fixed and Floating-Point Hardware Implementation for Enhanced Inverter Control: Leveraging FPGA in Recurrent Neural Network Applications View paper
- [14] Rigorous floating-point mixed-precision tuning View paper
- [15] Fooling a complete neural network verifier View paper
- [16] Scalable Neural Network Geometric Robustness Validation via HÃ¶lder Optimisation View paper
- [17] Trustworthy Machine Learning for High Assurance Systems View paper
- [18] Verifying Low-dimensional Input Neural Networks via Input Quantization View paper
- [19] Sound Floating-Point Neural Network Verification with MILP View paper
- [20] Nondeterminism-Aware Optimistic Verification for Floating-Point Neural Networks View paper
- [21] Sound Analysis of Floating-Point Programs View paper
- [22] Floating-Point Neural Network Verification at the Software Level View paper
- [23] Scaling up the static analysis of neural networks using affine forms View paper
- [24] Quantization with Guaranteed Floating-Point Neural Network Classifications View paper
- [25] A Certified Proof Checker for Deep Neural Network Verification in Imandra View paper
- [26] Taming Floating-Point Rounding Errors with Proofs (Invited Talk) View paper
- [27] QNNVerifier: A Tool for Verifying Neural Networks using SMT-Based Model Checking View paper
- [28] Certified Deep Learning: Verification and Training View paper
- [29] Correct Approximation of IEEE 754 Floating-Point Arithmetic for Program Verification View paper
- [30] Taxonomy and Techniques: Formal Verification of Adversarial Robustness in Deep Learning View paper
- [31] Probabilistic verification of neural networks using branch and bound View paper
- [32] Branch and Bound for Piecewise Linear Neural Network Verification View paper

- [33] Robustness verification of neural networks using polynomial optimization View paper
- [34] General Cutting Planes for Bound-Propagation-Based Neural Network Verification View paper
- [35] Bound Tightening Using Rolling-Horizon Decomposition for Neural Network Verification View paper
- [36] Quantitative verification of neural networks and its security applications View paper
- [37] A sound abstraction method towards efficient neural networks verification View paper
- [38] Exponent Relaxation of Polynomial Zonotopes and Its Applications in Formal Neural Network Verification View paper
- [39] Bridging neural ode and resnet: A formal error bound for safety verification View paper
- [40] Residual-based error bound for physics-informed neural networks View paper
- [41] Dynamic Back-Substitution in Bound-Propagation-Based Neural Network Verification View paper
- [42] Formal security analysis of neural networks using symbolic intervals View paper
- [43] Efficient neural network verification via adaptive refinement and adversarial search View paper
- [44] Verification of Neural Control Barrier Functions with Symbolic Derivative Bounds Propagation View paper
- [45] Symbolic-numeric programming in scientific computing View paper
- [46] Optimized Symbolic Interval Propagation for Neural Network Verification View paper
- [47] Achieving Verified Robustness to Symbol Substitutions via Interval Bound Propagation View paper
- [48] Formal techniques for verification and testing of cyber-physical systems View paper
- [49] Eager Falsification for Accelerating Robustness Verification of Deep Neural Networks View paper
- [50] Robust Training of Neural Networks against Bias Field Perturbations View paper
- [51] Framework design of Network intrusion detection based on convolutional neural networks View paper
- [52] Generating and checking dnn verification proofs View paper
- [53] Global optimization of objective functions represented by ReLU networks View paper
- [54] OSIP: Tightened Bound Propagation for the Verification of ReLU Neural Networks View paper
- [55] Improved geometric path enumeration for verifying relu neural networks View paper
- [56] Seev: Synthesis with efficient exact verification for relu neural barrier functions View paper
- [57] An abstract domain for certifying neural networks View paper
- [58] Complexity of Injectivity and Verification of ReLU Neural Networks (Extended Abstract) View paper