Novelty Assessment Report

Paper: Parameterized Hardness of Zonotope Containment and Neural Network Verification

PDF URL: https://openreview.net/pdf?id=y8N45EEW05

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-05

Abstract

Neural networks with ReLU activations are a widely used model in machine learning. It is thus important to have a profound understanding of the properties of the functions computed by such networks. Recently, there has been increasing interest in the (parameterized) computational complexity of determining these properties. In this work, we close several gaps and resolve an open problem posted by Froese et al. [COLT '25] regarding the parameterized complexity of various problems related to network verification. In particular, we prove that deciding positivity (and thus surjectivity) of a function $f\colon\mathbb{R}^d\to\mathbb{R}$ computed by a 2-layer ReLU network is W[1]-hard when parameterized by D. This result also implies that zonotope (non-)containment is W[1]-hard with respect to D, a problem that is of independent interest in computational geometry, control theory, and robotics. Moreover, we show that (a) approximating the maximum within any multiplicative factor in 2-layer ReLU networks, (b) computing the L_p-Lipschitz constant for $p\in(0,\infty]$ in 2-layer networks, and (c) approximating the L_p-Lipschitz constant in 3-layer networks are all NP-hard and W[1]-hard with respect to D. Notably, our hardness results are the strongest known so far and imply that the naive enumeration-based methods for solving these fundamental problems are all essentially optimal under the Exponential Time Hypothesis.

Core Task Landscape

This paper addresses: **Parameterized Complexity of Neural Network Verification and Zonotope Containment**

A total of **15 papers** were analyzed and organized into a taxonomy with **9 categories**.

Taxonomy Overview

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

- **Computational Complexity and Hardness Results**
- **Set Representation Methods for Neural Networks**
- **Reachability and Safety Verification of Neural Control Systems**

Complete Taxonomy Tree

- Parameterized Complexity of Neural Network Verification and Zonotope Containment Survey Taxonomy
- Computational Complexity and Hardness Results
  - Parameterized Complexity of Network Properties ★ (2 papers)
  - [0] Parameterized Hardness of Zonotope Containment and Neural Network Verification (Anon et al., 2026) View paper
  - [10] Complexity of Injectivity and Verification of ReLU Neural Networks (Froese, 2024) View paper
  - Geometric Containment and Intersection Complexity (2 papers)
  - [9] On the difficulty of intersection checking with polynomial zonotopes (Yushen Huang, 2023) View paper
  - [15] Technical Report Column1 (Kelley, 2025) View paper
- Set Representation Methods for Neural Networks
  - Zonotope-Based Abstract Domains (2 papers)
  - [1] The octatope abstract domain for verification of neural networks (Stanley Bak, 2023) View paper
  - [8] Zonotope Domains for Lagrangian Neural Network Verification (Jordan, 2022) View paper
  - Polynomial Zonotope Optimization (1 papers)
  - [4] Exponent relaxation of polynomial zonotopes and its applications in formal neural network verification (Tobias Ladner, 2024) View paper
  - Hybrid Zonotope Representations and Applications (2 papers)
  - [7] Hybrid Zonotopes Exactly Represent ReLU Neural Networks (Joshua Ortiz, 2023) View paper
  - [14] Unified Framework for Neural Network Verification with Hybrid Zonotopes (Deng, 2025) View paper
- Reachability and Safety Verification of Neural Control Systems
  - Hybrid Zonotope Reachability Analysis (3 papers)
  - [2] Reachability analysis using hybrid zonotopes and functional decomposition (Jacob A. Siefert, 2025) View paper
  - [5] Reachability Analysis of Neural Network Control Systems With Tunable Accuracy and Efficiency (Yuhao Zhang, 2024) View paper
  - [13] Reachability Analysis and Safety Verification of Neural Feedback Systems via Hybrid Zonotopes (Yuhao Zhang, 2022) View paper
  - Constrained and Symbolic Zonotope Verification (2 papers)
  - [11] Verification of Neural Network Control Systems using Symbolic Zonotopes and Polynotopes (Trapiello, 2023) View paper
  - [12] Safety Verification of Neural Feedback Systems Based on Constrained Zonotopes (Yuhao Zhang, 2022) View paper
  - Convolutional Network Reachability (1 papers)
  - [6] Efficient Reachability Analysis for Convolutional Neural Networks Using Hybrid Zonotopes (Zhang Yuhao, 2025) View paper
  - Provably-Safe Neural Network Training (1 papers)

◦ [3] Provably-Safe Neural Network Training Using Hybrid Zonotope Reachability Analysis (Kousik, 2025) View paper

Narrative

Core task: parameterized complexity of neural network verification and zonotope containment. The field organizes around three main branches. The first, Computational Complexity and Hardness Results, investigates fundamental limits—such as whether certain verification problems remain intractable even when restricted by natural parameters like network depth or width. The second branch, Set Representation Methods for Neural Networks, develops geometric abstractions (zonotopes, polynomial zonotopes, hybrid zonotopes, and related domains) that enable tractable overapproximation of reachable sets through neural layers. The third branch, Reachability and Safety Verification of Neural Control Systems, applies these representations to closed-loop settings where neural controllers interact with dynamical plants, aiming to certify safety properties end-to-end. Together, these branches reflect a progression from theoretical hardness insights through algorithmic tooling to practical verification workflows.

Within the complexity branch, a small handful of works probe how parameterizations affect tractability; for instance, ReLU Injectivity Complexity[10] examines structural properties of ReLU networks under dimension constraints, while Zonotope Containment Hardness[0] establishes lower bounds for deciding containment queries even when zonotope complexity is bounded. This line contrasts sharply with the representation-methods branch, where many studies refine zonotope variants—Hybrid Zonotopes Reachability[2], Safe Training Hybrid Zonotopes[3], Polynomial Zonotope Relaxation[4], and CNN Hybrid Zonotopes[6]—to balance precision and scalability. The original paper sits squarely in the hardness cluster, complementing ReLU Injectivity Complexity[10] by showing that zonotope containment remains computationally hard under natural parameterizations. Its emphasis on lower bounds provides a counterpoint to the algorithmic optimism in the representation branch, clarifying which problem features genuinely ease verification and which do not.

Related Works in Same Category

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

1. Complexity of Injectivity and Verification of ReLU Neural Networks

Authors: Froese, Vincent, Vincent Froese, Skutella, Martin, et al. (7 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Neural networks with ReLU activation play a key role in modern machine learning. Understanding the functions represented by ReLU networks is a major topic in current research as this enables a better interpretability of learning processes. Injectivity of a function computed by a ReLU network, that is, the question if different inputs to the network always lead to different outputs, plays a crucial role whenever invertibility of the function is required, such as, e.g., for inverse problems or gen...

Relationship Analysis

Both papers belong to the same taxonomy category analyzing parameterized complexity of fundamental neural network properties like positivity, injectivity, and surjectivity. They share overlapping focus on computational hardness results for ReLU network verification problems and connections to zonotope containment, with both establishing W[1]-hardness results parameterized by input dimension. The key difference is that the original paper provides more comprehensive hardness results including Lipschitz constant computation, approximation hardness, and stronger ETH-based lower bounds (nΩ(d) runtime), while the candidate paper focuses more narrowly on injectivity complexity (proving coNP-completeness) and presents positive algorithmic results (FPT algorithms for single-layer networks with small input dimension).

Contributions Analysis

Overall novelty summary. The paper establishes W[1]-hardness results for positivity and surjectivity of 2-layer ReLU networks, zonotope containment, and Lipschitz constant computation. It resides in the 'Parameterized Complexity of Network Properties' leaf, which contains only two papers total. This leaf sits within the broader 'Computational Complexity and Hardness Results' branch, a relatively sparse area compared to the representation-methods and reachability branches that dominate the taxonomy with over ten papers. The work thus occupies a niche focused on fundamental intractability rather than algorithmic development.

The taxonomy reveals that most neighboring research concentrates on set-based verification methods—zonotope variants, hybrid zonotopes, and polynomial zonotopes—designed to make verification tractable in practice. The sibling paper in the same leaf examines ReLU injectivity complexity, while the adjacent 'Geometric Containment and Intersection Complexity' leaf addresses zonotope containment from a geometric perspective. The original paper bridges these areas by proving that zonotope containment inherits hardness from network verification, connecting theoretical limits across both domains and contrasting sharply with the algorithmic optimism in neighboring branches.

Among 21 candidates examined, two contributions show potential prior overlap. The zonotope containment hardness claim (10 candidates examined, 1 refutable) and the Lipschitz constant hardness results (10 candidates examined, 1 refutable) each encounter one candidate that may provide overlapping prior work. The positivity/surjectivity hardness for 2-layer networks (1 candidate examined, 0 refutable) appears more novel within this limited search scope. The statistics suggest that while some hardness results may have precedent, the specific parameterized complexity angle and the connection to zonotope containment represent less-explored territory.

Given the limited search scope of 21 candidates from top-K semantic matches, the analysis captures immediate neighbors but cannot claim exhaustive coverage. The sparse population of the complexity-hardness branch and the paper's bridging role between network verification and geometric containment suggest meaningful contributions, though the refutable pairs indicate that certain hardness results may build incrementally on known lower bounds. A broader literature search would clarify whether the parameterized perspective and the zonotope connection constitute substantial novelty or refinements of established themes.

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

Contribution 1: W[1]-hardness of 2-layer ReLU network positivity and surjectivity

Description: The authors establish that determining whether a 2-layer ReLU network outputs a positive value (positivity) or is surjective is W[1]-hard with respect to the input dimension d, resolving an open problem from COLT '25. This result implies that simple enumeration algorithms are essentially optimal under the Exponential Time Hypothesis.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Training Fully Connected Neural Networks is -Complete

URL: View paper

Brief Assessment

Training Fully Connected Completeness[16] focuses on ∃R-completeness of training two-layer ReLU networks (empirical risk minimization), not on the W[1]-hardness of deciding positivity or surjectivity properties of already-trained networks. These are fundamentally different computational problems.

Contribution 2: W[1]-hardness of zonotope containment

Description: The authors prove that deciding whether one zonotope is contained in another is W[1]-hard when parameterized by dimension d. This follows from the duality between 2-layer ReLU networks and zonotopes, and has implications for computational geometry and robotics applications.

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. Complexity of Injectivity and Verification of ReLU Neural Networks

URL: View paper

Brief Assessment

ReLU Injectivity Complexity[10] focuses on injectivity and verification problems for ReLU networks, establishing coNP-completeness results. While it mentions zonotope containment in connection to surjectivity, it does not address the parameterized complexity (W[1]-hardness) of zonotope containment with respect to dimension d, which is the novel contribution claimed by the original paper.

---

2. On the co-NP-completeness of the zonotope containment problem

URL: View paper

Brief Assessment

Zonotope Containment Completeness[31] focuses on co-NP-completeness of zonotope containment, not parameterized complexity. No full text provided for detailed comparison.

---

3. Zonotope-Based Elastic Tube Model Predictive Control

URL: View paper

Brief Assessment

Elastic Tube MPC[28] focuses on tube-based model predictive control using zonotopic sets for robust control applications, not on the computational complexity or parameterized hardness of zonotope containment problems in computational geometry.

---

4. Combinatorial Optimization

URL: View paper

Brief Assessment

Combinatorial Optimization[29] is a workshop report covering various optimization topics. It does not address parameterized complexity of zonotope containment or provide technical results that would refute the original paper's novelty claim about W[1]-hardness.

---

5. Robust explicit model predictive control for hybrid linear systems with parameter uncertainties

URL: View paper

Brief Assessment

Robust Hybrid MPC[30] focuses on robust model predictive control using zonotope propagation for hybrid linear systems with parameter uncertainties, not on the parameterized computational complexity of zonotope containment problems.

---

6. Efficient computation of reachable sets of linear time-invariant systems with inputs

URL: View paper

Brief Assessment

Reachable Sets Computation[33] focuses on efficient algorithms for computing reachable sets of linear time-invariant systems using zonotopes, not on the parameterized complexity or hardness of zonotope containment problems.

---

7. Open Problem: Fixed-Parameter Tractability of Zonotope Problems

URL: View paper

Brief Assessment

Zonotope Tractability Problem[27] poses the open question of whether zonotope containment is fixed-parameter tractable with respect to dimension d, but does not claim to prove W[1]-hardness. The candidate explicitly states this complexity status remains 'unknown' and 'open'.

---

8. Resilient set-based state estimation for linear time-invariant systems using zonotopes

URL: View paper

Brief Assessment

Resilient Zonotope Estimation[32] focuses on set-based state estimation for LTI systems under sensor attacks using zonotopes for computational efficiency. It does not address the parameterized complexity or W[1]-hardness of zonotope containment problems in computational geometry.

---

9. Provably-Safe Neural Network Training Using Hybrid Zonotope Reachability Analysis

URL: View paper

Brief Assessment

Safe Training Hybrid Zonotopes[3] focuses on neural network training methods using hybrid zonotopes for reachability analysis and safety verification, not on the parameterized complexity of zonotope containment problems in computational geometry.

---

10. On the difficulty of intersection checking with polynomial zonotopes

URL: View paper

Prior Art Analysis

Polynomial Zonotope Intersection[9] demonstrates that prior work exists on the computational hardness of zonotope problems. The candidate paper proves that polynomial zonotope intersection checking is NP-hard and establishes complexity results for zonotope containment problems. While the original paper proves W[1]-hardness parameterized by dimension d, the candidate establishes NP-hardness for intersection checking with polynomial zonotopes, which is closely related to zonotope containment. The candidate's Corollary 1 explicitly states that 'polynomial zonotope intersection checking is np-hard,' and the paper discusses zonotope containment as a fundamental operation throughout. This demonstrates that computational hardness results for zonotope-related problems were established in prior work, though the specific parameterized complexity framework differs.

Evidence

Evidence 1 - **Rationale**: Both papers identify zonotope containment/intersection as a fundamental problem in robotics and control applications, with the candidate establishing hardness results for this operation prior to the original paper's publication. - **Original**: the latter asks whether one zonotope is contained within another, a question that has been extensively studied due to its applications in areas such as robotics and control - **Candidate**: determining whether the reachable set, represented as a polynomial zonotope, intersects an unsafe set is not straightforward. in fact, we show that this fundamental operation is np-hard, even for a simple class of polynomial zonotopes.

Evidence 2 - **Rationale**: The candidate paper discusses algorithms for zonotope intersection checking and their computational properties, establishing prior work on the algorithmic complexity of zonotope operations before the original paper's claims. - **Original**: zonotope (non-)containment is w[1]-hard with respect to d, a problem that is of independent interest in computational geometry, control theory, and robotics. - **Candidate**: setting aside the worst-case time complexity, the existing algorithm proposed to check for intersections, as well as perform plotting, is based on a combination of overapproximation using zonotopes and refinement using splitting [15,4]. is this algorithm guaranteed to converge to the true set, even g...

Contribution 3: W[1]-hardness results for Lipschitz constant computation and approximation

Description: The authors establish W[1]-hardness (parameterized by input dimension d) for computing and approximating Lipschitz constants in ReLU networks. They extend previous NP-hardness results to all p-norms and prove that approximation within any multiplicative factor remains hard, showing that enumeration-based methods are essentially optimal.

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. Semialgebraic Optimization for Lipschitz Constants of ReLU Networks

URL: View paper

Brief Assessment

Semialgebraic Lipschitz Optimization[26] focuses on semialgebraic optimization methods for computing Lipschitz constants, not on parameterized complexity theory or W[1]-hardness results. The technical approaches are fundamentally different.

---

2. LipBaB: Computing exact Lipschitz constant of ReLU networks

URL: View paper

Brief Assessment

LipBaB Exact Lipschitz[21] focuses on computing exact Lipschitz constants using branch-and-bound methods, not on complexity-theoretic hardness results. The candidate does not address parameterized complexity or W[1]-hardness.

---

3. Optimal lower Lipschitz bounds for ReLU layers, saturation, and phase retrieval

URL: View paper

Brief Assessment

Optimal Lower Lipschitz[25] focuses on deriving optimal lower Lipschitz bounds for ReLU layers using frame theory, not on computational complexity or hardness results for computing these constants.

---

4. Exactly Computing the Local Lipschitz Constant of ReLU Networks

URL: View paper

Prior Art Analysis

Local Lipschitz ReLU[17] establishes strong inapproximability results for computing Lipschitz constants of ReLU networks that predate the original paper's W[1]-hardness claims. Specifically, Local Lipschitz ReLU[17] proves that approximating the Lipschitz constant within any multiplicative factor that scales almost linearly with input dimension is hard under the Exponential Time Hypothesis (ETH), and extends NP-hardness results to $\ell 1$ and $\ell\infty$ norms. While the original paper claims to extend previous NP-hardness results to all p-norms and prove W[1]-hardness parameterized by input dimension d, Local Lipschitz ReLU[17] already demonstrated that 'it is provably hard to approximate the lipschitz constant of a network to within a factor that scales almost linearly with input dimension' and showed 'strong inapproximability results for estimating lipschitz constants of relu networks.' This prior work establishes the fundamental hardness of both exact computation and approximation of Lipschitz constants, challenging the novelty of the original paper's W[1]-hardness contribution.

Evidence

Evidence 1 - **Rationale**: Both papers establish hardness of approximating Lipschitz constants. Local Lipschitz ReLU[17] proves inapproximability with a factor scaling almost linearly with input dimension, which is a strong hardness result that predates the original paper's W[1]-hardness claims for approximation. - **Original**: we show that (a) approximating the maximum within any multiplicative factor in 2-layer relu networks, (b) computing the lp-lipschitz constant for p ∈ (0, ∞] in 2-layer networks, and (c) approximating the lp-lipschitz constant in 3-layer networks are all np-hard and w[1]-hard with respect to d. - **Candidate**: we show that that it is provably hard to approximate the lipschitz constant of a network to within a factor that scales almost linearly with input dimension.

Evidence 2 - **Rationale**: Local Lipschitz ReLU[17] extends the NP-hardness result from $\ell 2$ to other norms and establishes strong inapproximability by relating to maximum independent set, one of the hardest problems to approximate. This demonstrates prior work on extending hardness results beyond the original $\ell 2$ case. - **Original**: virmaux & scaman (2018) proved that computing thel2-lipschitz constant of a 2-layer relu network is np-hard. in section 6, we extend this to np-hardness for every p ∈ (0, ∞] and even w[1]-hardness with respect to d. - **Candidate**: exactly computing l2(f,rd) is np-hard [13]. this does not address the question of whether efficient approximation algorithms exist. we relate this problem to the problem of approximating the maximum independent set of a graph. maximum independent set is one of the hardest problems to approximate

---

5. Refining the inverse Lipschitz constant for injective ReLU networks

URL: View paper

Brief Assessment

Inverse Lipschitz Injective[19] focuses on inverse Lipschitz constants for injective ReLU layers, studying tightness of lower bounds. The original paper addresses forward Lipschitz constant computation complexity, which is a fundamentally different problem.

---

6. Near-optimal estimates for the -Lipschitz constants of deep random ReLU neural networks

URL: View paper

Brief Assessment

Random ReLU Lipschitz[18] studies Lipschitz constants of random ReLU networks with probabilistic bounds, not computational complexity. The candidate focuses on high-probability estimates for randomly initialized networks rather than worst-case hardness results for computing Lipschitz constants of arbitrary networks.

---

7. Upper and lower bounds for the Lipschitz constant of random neural networks

URL: View paper

Brief Assessment

Random Network Lipschitz Bounds[20] focuses on probabilistic upper and lower bounds for Lipschitz constants in random ReLU networks with specific weight initializations, not on computational complexity or hardness results for computing these constants.

8. Optimal robust memorization with relu neural networks

URL: View paper

Brief Assessment

Optimal Robust Memorization[24] focuses on constructing robust memorization networks and controlling Lipschitz constants for memorization tasks, not on the computational complexity of computing or approximating Lipschitz constants in general ReLU networks. The paper does not address W[1]-hardness or parameterized complexity of Lipschitz constant computation.

9. Analytical Bounds on the Local Lipschitz Constants of ReLU Networks

URL: View paper

Brief Assessment

Analytical Lipschitz Bounds[22] focuses on deriving analytical upper bounds for local Lipschitz constants of ReLU networks, not on computational complexity or hardness results. The candidate does not address W[1]-hardness, parameterized complexity, or approximation hardness for computing Lipschitz constants.

10. Three quantization regimes for ReLU networks

URL: View paper

Brief Assessment

Three Quantization Regimes[23] focuses on approximation of Lipschitz functions by quantized ReLU networks, not on the computational complexity of computing Lipschitz constants. The paper does not address W[1]-hardness or parameterized complexity theory.

Appendix: Text Similarity Detection

Textual similarity detection checked 21 papers and found 2 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Complexity of Injectivity and Verification of ReLU Neural Networks

Detected in: Core Task (sibling), Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Parameterized Hardness of Zonotope Containment and Neural Network Verification View paper
- [1] The octatope abstract domain for verification of neural networks View paper
- [2] Reachability analysis using hybrid zonotopes and functional decomposition View paper
- [3] Provably-Safe Neural Network Training Using Hybrid Zonotope Reachability Analysis View paper
- [4] Exponent relaxation of polynomial zonotopes and its applications in formal neural network verification View paper
- [5] Reachability Analysis of Neural Network Control Systems With Tunable Accuracy and Efficiency View paper
- [6] Efficient Reachability Analysis for Convolutional Neural Networks Using Hybrid Zonotopes View paper
- [7] Hybrid Zonotopes Exactly Represent ReLU Neural Networks View paper
- [8] Zonotope Domains for Lagrangian Neural Network Verification View paper
- [9] On the difficulty of intersection checking with polynomial zonotopes View paper
- [10] Complexity of Injectivity and Verification of ReLU Neural Networks View paper
- [11] Verification of Neural Network Control Systems using Symbolic Zonotopes and Polynotopes View paper
- [12] Safety Verification of Neural Feedback Systems Based on Constrained Zonotopes View paper
- [13] Reachability Analysis and Safety Verification of Neural Feedback Systems via Hybrid Zonotopes View paper
- [14] Unified Framework for Neural Network Verification with Hybrid Zonotopes View paper
- [15] Technical Report Column1 View paper
- [16] Training Fully Connected Neural Networks is -Complete View paper
- [17] Exactly Computing the Local Lipschitz Constant of ReLU Networks View paper
- [18] Near-optimal estimates for the -Lipschitz constants of deep random ReLU neural networks View paper
- [19] Refining the inverse Lipschitz constant for injective ReLU networks View paper
- [20] Upper and lower bounds for the Lipschitz constant of random neural networks View paper
- [21] LipBaB: Computing exact Lipschitz constant of ReLU networks View paper
- [22] Analytical Bounds on the Local Lipschitz Constants of ReLU Networks View paper
- [23] Three quantization regimes for ReLU networks View paper
- [24] Optimal robust memorization with relu neural networks View paper
- [25] Optimal lower Lipschitz bounds for ReLU layers, saturation, and phase retrieval View paper
- [26] Semialgebraic Optimization for Lipschitz Constants of ReLU Networks View paper
- [27] Open Problem: Fixed-Parameter Tractability of Zonotope Problems View paper
- [28] Zonotope-Based Elastic Tube Model Predictive Control View paper
- [29] Combinatorial Optimization View paper
- [30] Robust explicit model predictive control for hybrid linear systems with parameter uncertainties View paper
- [31] On the co-NP-completeness of the zonotope containment problem View paper
- [32] Resilient set-based state estimation for linear time-invariant systems using zonotopes View paper
- [33] Efficient computation of reachable sets of linear time-invariant systems with inputs View paper