Reducing Overestimation in Interval Neural Networks Using Certified Polynomial Approximations and Partitioning for Trustworthy Clinical AI

María L. Reyna-Cruz, Kristalys Ruiz-Rohena, Isaiah D. Johnson,
Vladik Kreinovich, Martine Ceberio, and Christoph Lauter

Abstract Reliable uncertainty estimation is essential for trustworthy artificial intelligence, particularly in safety-critical applications such as medical imaging and clinical decision support. Interval neural networks provide mathematically guaranteed output bounds by propagating interval-valued inputs through neural network layers. However, interval propagation often suffers from severe overestimation, especially when evaluating nonlinear activation functions, causing uncertainty bounds to widen rapidly as information propagates through the network.

This work investigates certified polynomial approximations and partitioning as a mechanism for reducing interval overestimation. Given an activation function over a bounded domain, minimax polynomial approximations are generated using the `Sollya` framework together with certified uniform error bounds. These approximations produce guaranteed interval enclosures while avoiding direct interval eval-

María L. Reyna-Cruz
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: mlchamorro@miners.utep.edu

Kristalys Ruiz-Rohena
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: kruizrohena@miners.utep.edu

Isaiah D. Johnson
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: idjohnson2@miners.utep.edu

Vladik Kreinovich
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: vladik@utep.edu

Martine Ceberio
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: mceberio@utep.edu

Christoph Lauter
Department of Computer Science, University of Texas at El Paso, 500 W. University
El Paso, Texas 79968, USA, e-mail: cqlauter@utep.edu

uation of nonlinear activation functions.
To further improve enclosure tightness, the input domain is partitioned into local subintervals and independent certified polynomial approximations are constructed for each partition. Preliminary experiments on the tanh and sigmoid activation functions demonstrate substantially tighter interval enclosures than global polynomial approximations while preserving mathematical guarantees.
Although the present study focuses on activation functions, the proposed methodology constitutes a fundamental building block toward reliable uncertainty propagation in interval neural networks for trustworthy clinical artificial intelligence.

1 Introduction

Artificial intelligence has become increasingly integrated into healthcare applications including medical image analysis, disease diagnosis, prognosis, and clinical decision support. While modern deep learning models often achieve remarkable predictive performance, they frequently provide overconfident prediction. This limitation has motivated growing interest in trustworthy artificial intelligence methods capable of quantifying predictive uncertainty alongside model predictions.

One mathematically rigorous framework for uncertainty representation is interval analysis. Rather than propagating single numerical values, interval methods propagate sets of possible values through every computational operation, producing guaranteed enclosures of all feasible outputs. Such guarantees are particularly attractive for safety-critical domains, where understanding the range of plausible predictions may be as important as the prediction itself.

Clinical decision-making naturally relies on reasoning under uncertainty. Physicians rarely interpret laboratory values, imaging measurements, or physiological variables as exact quantities; instead, they reason within acceptable ranges while accounting for measurement variability, observer dependence, and biological heterogeneity. In medical imaging, for example, anatomical landmarks may vary between operators, image quality may be suboptimal, and patient-specific factors introduce additional uncertainty. Consequently, AI systems intended to support clinical decisions should not only provide predictions but also communicate how reliable those predictions are. Mathematical frameworks capable of propagating uncertainty while preserving rigorous guarantees therefore represent an attractive direction for trustworthy clinical AI.

Interval neural networks extend this principle to deep learning by representing neuron activations as intervals and propagating uncertainty throughout successive network layers. Unfortunately, interval arithmetic is well known to suffer from overestimation [10]. Dependency effects and repeated nonlinear operations progressively enlarge interval widths, often producing uncertainty bounds that become too conservative to remain practically informative.

Among the principal contributors to this phenomenon are nonlinear activation functions. Smooth nonlinearities such as tanh and sigmoid introduce curvature that significantly amplifies interval widening. Even for moderate input domains, direct interval evaluation can produce enclosures substantially wider than the true function range. Since activation functions are evaluated throughout every hidden layer, these local overestimations accumulate and may eventually dominate the uncertainty propagated by the entire network.

Several approaches have been proposed to approximate nonlinear functions while preserving mathematical guarantees. Polynomial approximations are particularly attractive because polynomial evaluation behaves considerably better under interval arithmetic than direct evaluation of transcendental functions [9]. Moreover, minimax approximations computed in the supremum norm provide nearly optimal uniform approximations over bounded intervals.

In this work we investigate certified polynomial approximations generated with the `Sollya` framework [4, 12]. For each activation function, `Sollya` computes minimax polynomial approximations using the Remez algorithm [1] together with certified uniform approximation errors obtained through the `dirtyinfnorm` algorithm [3]. These certified error bounds guarantee that the true function remains enclosed within the polynomial approximation expanded by a rigorously computed error interval [5].

Beyond global polynomial approximation, we propose partitioning the approximation domain into multiple subintervals and constructing independent certified approximations over each partition. Because each local polynomial is optimized over a smaller domain, approximation errors decrease substantially, yielding tighter certified interval enclosures while preserving rigorous correctness guarantees.

Although this study focuses on activation functions rather than complete neural network architectures, activation evaluation represents one of the primary sources of interval overestimation during forward propagation. Consequently, improving activation enclosures constitutes an important step toward scalable interval neural networks capable of providing reliable uncertainty estimates for medical imaging, computer-aided diagnosis, and other trustworthy clinical AI applications.

The main contributions of this work are:

- formulation of certified minimax polynomial approximations for interval evaluation of smooth activation functions;
- a partitioning strategy that further tightens certified interval enclosures through local polynomial approximation;
- experimental evaluation on tanh and sigmoid activation functions demonstrating substantial reductions in interval overestimation;
- discussion of how these techniques can serve as building blocks for future interval neural networks applied to trustworthy medical artificial intelligence.

2 Background & Related Work

Interval arithmetic provides mathematically rigorous enclosures for numerical computations and has been extensively studied for reliable scientific computing and uncertainty propagation [10]. Although interval methods guarantee that the exact solution remains enclosed, repeated interval operations often introduce dependency effects that produce significant overestimation.

Several approaches have been proposed to reduce interval overestimation, including affine arithmetic, Taylor models, constraint propagation, and polynomial approximations. Among these alternatives, certified polynomial approximations are interesting because polynomial evaluation is computationally efficient while maintaining rigorous mathematical guarantees.

The `Sollya` framework provides a mature environment for constructing minimax polynomial approximations of elementary functions and has served as one of the foundations for automatic mathematical function code generators such as Metalibm [4, 7]. Polynomial coefficients are computed using implementations of the Remez algorithm [1], while certified approximation errors are estimated using the `dirtyinfnorm` algorithm [3]. These capabilities have made `Sollya` an important component in the development of correctly rounded mathematical libraries and numerical software.

Domain splitting has previously been investigated for improving polynomial approximations by constructing independent approximations over smaller intervals rather than using a single global polynomial [8]. Since local approximations are computed over reduced domains, they typically achieve smaller approximation errors than a single global polynomial of the same degree.

Although interval neural networks have received increasing attention for uncertainty-aware machine learning, relatively little work has focused on reducing overestimation introduced by nonlinear activation functions through certified approximation techniques. The present work investigates this problem and studies how certified local polynomial approximations may serve as building blocks for reliable interval neural network propagation.

3 Theoretical Backbone

3.1 Baseline: Interval Overestimation

Let $f : [a,b] \to \mathbb{R}$ be a function. If f is continuous for all $x \in [a,b]$, where $x \in \mathbb{R}$, then a standard result from interval analysis states that the exact range of f over $[a,b]$ is given by [10]

$$f([a,b]) = \left[\inf_{x\in[a,b]} f(x), \ \sup_{x\in[a,b]} f(x) \right].$$

In practice, interval arithmetic produces an enclosure that contains this range, but may be wider due to dependency effects and nonlinearities. This gap is referred to as *overestimation*.

When interval computations are repeatedly applied throughout multiple neural-network layers, this widening accumulates, producing uncertainty bounds that may become substantially larger than the true output uncertainty.

3.2 Polynomial-Based Interval Enclosure

Let $f : [a,b] \to \mathbb{R}$ be a function and $P_n(x)$ be a polynomial such that

$$|f(x) - P_n(x)| \leq \varepsilon \quad \text{for all } x \in [a,b].$$

By standard interval enclosure arguments [10], the range of f over $[a,b]$ satisfies

$$f([a,b]) \subseteq [\min P_n([a,b]) - \varepsilon,\ \max P_n([a,b]) + \varepsilon]\,.$$

Instead of directly evaluating the nonlinear activation function under interval arithmetic, we evaluate its polynomial approximation and expand the resulting interval by by a known error bound. The resulting interval remains guaranteed to contain the true function values while often being substantially tighter than direct interval evaluation.

3.3 Partitioned Approximation

This partitioning strategy follows the idea by Kupriianova and Lauter [8] of piecewise polynomial approximation, where dividing the domain into subintervals reduces approximation error compared to global polynomial approximations as proposed by Burden and Faires [2].

We partition the interval $[a,b]$ into smaller subintervals (e.g., $[a_1,b_1]$, $[a_2,b_2]$, $[a_3,b_3]$).

For each subinterval, we compute a polynomial approximation $p_i(x)$ with an associated error bound ε_i, such that

$$|f(x) - p_i(x)| \leq \varepsilon_i \quad \text{for all } x \in [a_i,b_i].$$

This yields a guaranteed enclosure on each subinterval:

$$f([a_i,b_i]) \subseteq [\inf p_i([a_i,b_i]) - \varepsilon_i,\ \sup p_i([a_i,b_i]) + \varepsilon_i]\,.$$

The overall bound over $[a,b]$ is obtained by taking the smallest interval that contains all subinterval bounds. In practice, this corresponds to taking the minimum of all lower bounds and the maximum of all upper bounds:

$$f([a,b]) \subseteq \left[\min_i \left(\inf p_i([a_i,b_i]) - \varepsilon_i\right),\ \max_i \left(\sup p_i([a_i,b_i]) + \varepsilon_i\right)\right].$$

Each subinterval provides a tight local bound, and the final enclosure is obtained by covering all these local bounds with a single interval.

4 Proposed Methodology

The goal of the proposed methodology is to reduce interval overestimation produced during the evaluation of nonlinear activation functions while preserving mathematically certified bounds. Rather than evaluating transcendental activation functions directly through interval arithmetic, we replace each activation evaluation with a certified polynomial enclosure. The proposed workflow consists of the following steps:

1. Select an activation function $f(x)$ and an input interval $[a,b]$.
2. Compute a global degree-5 minimax polynomial approximation using the `remez` algorithm.
3. Compute the certified uniform approximation error using the `dirtyinfnorm` routine.
4. Evaluate the polynomial over the interval and enlarge the resulting enclosure by the certified error.
5. Partition the input domain into three subintervals.
6. Repeat the polynomial approximation and error certification independently on each subinterval.
7. Merge the local certified enclosures into the final interval enclosure.
8. Compare the resulting interval width against the global polynomial enclosure and the true function range.

Although the present work evaluates activation functions independently, the proposed procedure is intended as a modular component of interval neural-network inference. During forward propagation, each nonlinear activation could be replaced by its corresponding certified polynomial enclosure, allowing uncertainty to propagate through the network with substantially reduced interval overestimation. This capability is particularly relevant for trustworthy medical AI, where reliable uncertainty estimates are desirable for clinical decision support and medical image analysis.

5 Results

We evaluate the proposed interval tightening approach on the tanh and sigmoid activation functions over the domain $[-2,2]$ with polynomials of degree 5. Degree-5 polynomials were selected as a compromise between approximation accuracy and polynomial complexity. Lower-degree approximations produced noticeably larger approximation errors, while higher-degree approximations did not significantly improve enclosure tightness for the considered intervals. Then, three subintervals were selected to capture the main curvature variations of the activation functions over $[-2, 2]$ while maintaining low partition complexity.

For tanh, a global degree-5 minimax polynomial approximation produces a highly conservative enclosure of approximately $[-4.306,\ 4.306]$, while the actual function range is $[-0.964,\ 0.964]$. By partitioning the domain into three subintervals and constructing local minimax polynomial approximations using `Sollya`, we obtain a significantly tighter certified enclosure of approximately $[-1.378,\ 1.378]$. The absolute approximation error remains below 4×10^{-4} across all subintervals and reaches values as low as 10^{-6} in flatter regions.

Additionally, we evaluate the sigmoid activation function. A global degree-5 approximation produces an enclosure of approximately $[-0.189,\ 1.189]$, while the actual function range is $[0.119,\ 0.881]$. Using the same partitioning strategy yields a substantially tighter enclosure of approximately $[0.066,\ 0.934]$.

These experiments show that local certified polynomial approximations significantly reduce interval overestimation for activation functions.

Table 1 Comparison of global and partitioned polynomial enclosures.

Function	Method	Enclosure	True Range
tanh	Global deg-5	$[-4.306,\ 4.306]$	$[-0.964,\ 0.964]$
	Partitioned	$[-1.378,\ 1.378]$	
Sigmoid	Global deg-5	$[-0.189,\ 1.189]$	$[0.119,\ 0.881]$
	Partitioned	$[0.066,\ 0.934]$	

6 Discussion

The experiments demonstrate that certified polynomial approximations substantially reduce interval overestimation while preserving guaranteed bounds. Local approximations over subintervals improve enclosure tightness by adapting to variations in function curvature.

Approximation quality varies across the domain. Regions near the origin, where tanh exhibits stronger curvature, require larger error bounds than flatter outer regions, suggesting that adaptive partitioning strategies may further improve performance.

Preliminary experiments with non-smooth activations such as ReLU and ELU revealed limitations of direct minimax approximation using `Sollya`. Because these activations are piecewise-defined and not globally smooth, polynomial approximations become less stable near transition points.

Floating-point rounding is also important in reliable numerical computation. Recent efforts such as the CORE-MATH project further illustrate the continuing importance of correctly rounded elementary functions for reliable scientific software

[11], although detailed rounding-error propagation analysis is outside the scope of the present work.

Overall, the proposed framework provides a theoretically grounded approach for reliable uncertainty propagation in interval neural networks.

References

1. N. Brisebarre and S. Chevillard, "Efficient polynomial L^∞-approximations", In: *Proceedings of the 18th IEEE Symposium on Computer Arithmetic ARITH'07*, Montpellier, France, June 25–27, 2007, 0-7695-2854-6/07
2. R. L. Burden and J. D. Faires, *Numerical Analysis*, 10th edition, Cengage Learning, 2015.
3. S. Chevillard, J. Harrison, M. Joldes, and C. Lauter, "Efficient and accurate computation of upper bounds of approximation errors", *Theoretical Computer Science*, 2011, Vol. 412, No. 16, pp. 1523–1543, doi 10.1016/j.tcs.2010.11.052
4. S. Chevillard, M. Joldeş, and C. Lauter, "Sollya: an environment for the development of numerical codes", In: *Mathematical Software – ICMS 2010*, Springer, Heidelberg, Germany, 2010, pp. 28–31.
5. M. Daumas and G. Melquiond, "Certification of bounds on expressions involving rounded operators", *ACM Transactions on Mathematical Software*, 2010, Vol. 37, No. 1, 2010, doi 10.1145/1644001.1644003
6. L. Fousse, G. Hanrot, V. Lefèvre, P. Pélissier, and P. Zimmermann, "MPFR: A Multiple-Precision Binary Floating-Point Library with Correct Rounding", *ACM Transactions on Mathematical Software*, 2007, Vol. 33, No. 2, Paper 13, doi 10.1145/1236463.1236468
7. O. Kupriianova and C. Lauter, "Metalibm: A mathematical functions code generator", In: H. Hongand and C. Yap (eds.), *Mathematical Software – ICMS 2014*, Springer, Berlin, Heidelberg, 2014, pp. 713–717.
8. O. Kupriianova and C. Lauter, "A domain splitting algorithm for the mathematical functions code generator", *Poroceedings of the 48th Asilomar Conference on Signals, Systems and Computers*, California, USA, November 2–5, 2014, doi 10.1109/ACSSC.2014.7094664
9. J.-M. Muller, *Elementary Functions: Algorithms and Implementation*, Springer, 2006.
10. J.-M. Muller, N. Brunie, F. Dinechin, C.-P. Jeannerod, M. Joldes, V. Lefèvre, G. Melquiond, N. Revol, and S. Torres, "Interval arithmetic", In: In: *Handbook of Floating-Point Arithmetic*, Springer, 2018, pp. 453–477, doi 10.1007/978-3-319-76526-6_12
11. A. Sibidanov, P. Zimmermann, and S. Glondu, "The core-math project", In: *Proceedings of the 29th IEEE Symposium on Computer Arithmetic ARITH'22*, September 12–14, 2022, pp. 26–34, doi 10.1109/ARITH54963.2022.00014
12. L. N. Trefethen, *Approximation Theory and Approximation Practice, Extended Edition*, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, USA, 2019, doi 10.1137/1.9781611975949