
ZONOGPT: TOWARDS AN ABSTRACT DOMAIN FOR VERIFYING LARGE GPT MODELS

Hai Duong

Department of Computer Science
George Mason University
Fairfax, VA, USA

Thanh Le

Unaffiliated
Yokosuka, Japan

ThanhVu Nguyen

Department of Computer Science
George Mason University
Fairfax, VA, USA

ABSTRACT

Transformer-based models are widely used for reasoning, coding, and multimodal agentic tasks. To provide formal assurance of desirable behaviors, such as robustness, safety, and fairness, neural network verification techniques prove required properties and provide auditable guarantees before deployment. However, prior work remains limited to small or restricted Transformers, and maintaining precision across deep models remains challenging. In this work, we introduce ZONOGPT, an abstract domain for verifying large transformers that maintains a space complexity independent of network depth. ZONOGPT uses a structured zonotope and a generator reduction mechanism to efficiently preserve correlations. To maintain precision, it introduces block-specific fused transformations for Attention and LayerNorm that retain feature relations, along with an affine transform for GELU that preserves generator relations. These mechanisms enable ZONOGPT to be the first approach to verify standard architectures, scaling to official HuggingFace models up to GPT-2 Medium (24 blocks, 300M+ parameters) and successfully verifying 1,339 instances across text and vision tasks.

1 Introduction

Transformer-based models [Vaswani et al., 2017], *e.g.*, generative pretrained transformer (GPT) [Radford et al., 2019], are widely used for reasoning, coding, and multimodal agentic tasks [Comanici et al., 2025, Liu et al., 2024]. However, adversarial prompts can change task predictions [Xu et al., 2024a] or bypass model safeguards [Zou et al., 2023, Chu et al., 2025]. Furthermore, compression and quantization allow open-weight models to run on consumer hardware without cloud infrastructure [Yang et al., 2025]. Local deployments can bypass cloud moderation and increase misuse risks [Huang et al., 2024].

As traditional software, empirical robustness evaluation is necessary but insufficient since it only exposes failures on tested inputs [Zhang et al., 2025, Joo et al., 2025]. To provide formal assurance that these models have desirable behaviors, such as robustness [Jia et al., 2019, Zhu et al., 2024], safety [Zou et al., 2023] and fairness [Gallegos et al., 2024, Wang et al., 2023], researchers have developed Neural Network Verification (NNV) techniques that prove required properties of neural networks, which include Transformer-based models, and provide auditable guarantees before release or deployment [Shi et al., 2020, Bonaert et al., 2021, Huang et al., 2026, Duong et al., 2025a].

While abstraction [Cousot and Cousot, 1977, Singh et al., 2018, Wang et al., 2021] has successfully scaled NNV to large models [Kaulen et al., 2025], verification remains limited to small or restricted transformers. Existing work handles at most three small simplified blocks [Shi et al., 2020, Huang et al., 2026]. Bonaert et al. [2021] scales to 12 blocks with much smaller configurations. None of these tools scale to standard transformer architectures, *e.g.*, GPT-2 [Radford et al., 2019]. We discuss these works in more detail in §B.

In addition, beyond scalability, maintaining precision across deep transformers remains challenging, *e.g.*, 24 blocks in GPT-2 Medium. Abstract domains inherently lose precision as they propagate through successive non-linear layers [Wang et al., 2018, Singh et al., 2019], *e.g.*, while CoVEnN scales via decomposition [Duong et al., 2025b], its interval bounds between blocks substantially degrade precision. Furthermore, recent techniques only tighten isolated operations, such as softmax [Wei et al., 2023] or attention products [Zhang et al., 2024, Huang et al., 2026], and does not address the over-approximation error that accumulates across layers.

In this paper, we introduce ZONOGPT, an abstract domain for verifying large transformers that maintains a space complexity (*e.g.*, memory footprint) independent of network depth while preserving relations between tokens and internal features. First, ZONOGPT uses structured zonotopes capturing: (i) global sequence correlations through shared generators and (ii) localized variations through local generators (§4.1). Second, reduction mechanisms limit number of generators after each block (§4.4), ensuring memory footprint remain bounded and independent of network depth (§4.5).

To maintain precision, ZONOGPT introduces block-specific abstract transformations tailored to transformer components. Particularly, ZONOGPT applies a fused transform for attention residual block (§4.2) that jointly over-approximates the skip and attention branches. By preserving the shared generator structure between the two, ZONOGPT yields a tighter bound than bounding each branch independently. For the MLP residual block, ZONOGPT further improves precision by employing a fused transform for LayerNorm to preserve feature relations within each token (§4.3.1), along with an affine transform for GELU (§4.3.2).

Evaluated on 2,880 text and vision instances, ZONOGPT successfully verifies 1,339 instances (46.5%), scaling to official HuggingFace models up to GPT-2 Medium (24 blocks, 300M+ parameters). As perturbation increases, fused attention proves critical for deep configurations (4+ blocks), *e.g.*, 867 verified instances (40.1%), while unfused variants verify fewer than 1%.

Contributions. The primary contributions of this paper include: (i) A structured zonotope (§4.1) and generator reduction mechanism (§4.4) that preserve global and local relations with a depth-independent space complexity; (ii) Fused transforms for attention (§4.2) and LayerNorm (§4.3.1) that preserve feature relations, and an affine GELU transform (§4.3.2) that retains input generators; (iii) Verification of official HuggingFace GPT-2 architectures adapted for vision tasks, scaling to GPT-2 Medium (300M+ parameters) under full-dimensional input perturbations (§5).

While the GPT models considered by ZONOGPT is still very small compared to the largest ones, it is the first approach to verify standard architectures, scaling to official HuggingFace models up to GPT-2 Medium (24 blocks, 300M+ parameters) (§5). We view this as a first step towards verifying larger GPT models, and we hope that our work will inspire future research in this direction.

2 Background

Neural Network Verification. Given a neural network N and a property ϕ , the NNV problem asks whether ϕ is a valid property of N . Typically, ϕ is a formula $\phi_{in} \Rightarrow \phi_{out}$, where ϕ_{in} and ϕ_{out} are properties over the inputs and outputs of N , respectively. An NNV tool searches for a *counterexample* input to N that satisfies ϕ_{in} but violates ϕ_{out} : if none exists, ϕ is valid; otherwise, ϕ is invalid [Wang et al., 2021, Ferrari et al., 2022, Wu et al., 2024, Duong et al., 2023, 2024, 2025c, Duong and Nguyen, 2025]. Modern verifiers often use branch-and-bound [Bunel et al., 2020, Duong et al., 2026a], whose bounding step computes all reachable outputs efficiently [Singh et al., 2018, 2019, Wang et al., 2018, Zhou et al., 2024, 2025].

Among those, zonotope [Singh et al., 2018] is the abstraction that balances precision and scalability:

$$\mathcal{Z}(c, G) = \{c + G\epsilon : \epsilon \in [-1, 1]^m\} \quad (1)$$

where c is the center and each column of G is a generator controlled by one coefficient in ϵ . The same coefficient can influence multiple outputs, allowing a zonotope to preserve their relations. Zonotope transformation for affine layer is exact, *e.g.*, $W\mathcal{Z}(c, G) + b = \mathcal{Z}(Wc + b, WG)$, while nonlinear operations require additional generators to bound their approximation errors. These generators improve precision but make the computational complexity grow with network sizes.

ZONOGPT builds on zonotopes with shared generators for cross-token and local generators for token-specific relations (§4.1). Generator reduction then limits their number after each block (§4.4), preserving the relations needed by Transformer operations without growth across network depth.

Residual Taylor Expansion. Let $A(x) = x + B(x)$ be a twice continuously differentiable residual, $Z = \{c + \sum_j H_j \epsilon_j : |\epsilon_j| \leq 1\}$ be its input zonotope. For every $x \in Z$, the change from the center is $v = x - c = \sum_j H_j \epsilon_j$. Define $g_v(t) = A(c + tv)$ for $t \in [0, 1]$. Taylor’s theorem along g_v gives

$$A(c + v) = A(c) + J_A(c)v + e_A(c, v), \quad e_A(c, v) = \int_0^1 (1 - t)g_v''(t) dt \quad (2)$$




Fig. 1: Running example for the attention residual for $S = 2$ tokens, hidden size $d = 2$, one head.

where $J_f(c)$ denotes the Jacobian matrix of a function f evaluated at c . Since $A(x) = x + B(x)$, the identity $J_A(c) = I + J_B(c)$ turns the linear term in Eq. 2 into $\sum_j (H_j + J_B(c)H_j)\epsilon_j$. Thus, the Taylor expansion of $A(Z)$ is:

$$A(Z) \subseteq Z' = \left\{ A(c) + \sum_j (H_j + J_B(c)H_j)\epsilon_j + e_A(c, v) : |\epsilon_j| \leq 1 \right\} \quad (3)$$

ZONOGPT uses Eq. 3 to construct the output zonotope for attention (§4.2) and LayerNorm (§4.3.1).

3 Motivating Example

We demonstrate ZONOGPT on an attention-residual block with $S = 2$ tokens and hidden size $d = 2$ (Fig. 1). We verify that its first output coordinate remains positive given an ℓ_∞ perturbation:

$$\varphi \equiv \|x - x_0\|_\infty \leq r \implies [A(x)]_{1,1} > 0, \quad x_0 = [[0.0, 0.0], [0.7, 0.8]], \quad r = 0.02 \quad (4)$$

Scalability. DEEPZ represents the input with Sd dense generators, each containing Sd coefficients, and therefore stores S^2d^2 coefficients at the input. For the running example ($S = 2, d = 2$):

$$\begin{aligned} x_{11} &= x_{11}^0 + r\epsilon_1 + 0\epsilon_2 + 0\epsilon_3 + 0\epsilon_4 & x_{12} &= x_{12}^0 + 0\epsilon_1 + r\epsilon_2 + 0\epsilon_3 + 0\epsilon_4 \\ x_{21} &= x_{21}^0 + 0\epsilon_1 + 0\epsilon_2 + r\epsilon_3 + 0\epsilon_4 & x_{22} &= x_{22}^0 + 0\epsilon_1 + 0\epsilon_2 + 0\epsilon_3 + r\epsilon_4 \end{aligned} \quad (5)$$

DEEPZ stores all 16 coefficients in Eq. 5, e.g., $(r, 0, 0, 0)$ for x_{11} , although each generator affects only one token. ZONOGPT stores the same input without the cross-token zeros:

$$x_{11} = x_{11}^0 + r\epsilon_1 + 0\epsilon_2 \quad x_{12} = x_{12}^0 + 0\epsilon_1 + r\epsilon_2 \quad x_{21} = x_{21}^0 + r\epsilon_3 + 0\epsilon_4 \quad x_{22} = x_{22}^0 + 0\epsilon_3 + r\epsilon_4 \quad (6)$$

ZONOGPT stores eight coefficients (instead of the 16). For the GPT-2-small input with $S = 512$ and $d = 768$, ZONOGPT uses $Sd^2 \cdot 8/2^{30} = 2.25$ GiB, while DEEPZ uses $S^2d^2 \cdot 8/2^{30} = 1152$ GiB. For hidden states after each residual component, ZONOGPT stores generators at the source token:

$$Z_s = c_s + \sum_{j=1}^m G_{j,s}\epsilon_j + L_s\epsilon_s^{\text{loc}} + [-b_s, b_s] \quad (7)$$

ZONOGPT uses G for shared generators, L_s for generators local to token s , and b_s for the remaining independent values. With at most m shared and q local generators per token, these terms require $O(Sd(m+q+1))$ coefficients independent of network depth (§4.5). In contrast, DEEPZ adds $O(Sd)$ dense generators per block. Each dense generator contains Sd coefficients, so one block adds $O(S^2d^2)$ coefficients and L blocks add $O(LS^2d^2)$.

Precision. We use the network in Fig. 1 with the following fixed parameters:

$$W_{qkv}^\top = \begin{bmatrix} 3.1 & -0.2 & 1.0 & 0.1 & 0.8 & -3.0 \\ 0.7 & -1.1 & -4.3 & 1.2 & -0.5 & 0.6 \end{bmatrix}, \quad W_o = \begin{bmatrix} 0.112 & 0.029 \\ 0.014 & 0.153 \end{bmatrix} \quad (8)$$

$b_{qkv} = (-0.1, 0.1, 0, 0.2, 0.2, -0.2)$, $b_o = (0.32122, 0.95628)$, $\gamma = (1.6, 1.8)$, $\beta = (0.2, 0.3)$, and $\tau = 10^{-5}$. We represent the input interval by the zonotope

$$X = \begin{bmatrix} 0 & 0 \\ 0.7 & -0.8 \end{bmatrix} + 0.02 \begin{bmatrix} \epsilon_{11} & \epsilon_{12} \\ \epsilon_{21} & \epsilon_{22} \end{bmatrix}, \quad c = \begin{bmatrix} 0 & 0 \\ 0.7 & -0.8 \end{bmatrix}, \quad \epsilon_{ij} \in [-1, 1] \quad (9)$$

The top row of Fig. 2 follows DEEPZ through LayerNorm, $QK^\top$, and softmax. Each nonlinear operator bounds its input with fresh generators and the attention output has growth to 34 generators:

$$[A(X)]_{1,1} = z_0 + \sum_{i=1}^{34} g_i\epsilon_i, \quad z_0 \approx 0.36, \quad \sum_i |g_i| \approx 1.61 \quad (10)$$




Fig. 2: First-token projections for operator-wise DEEPZ (top) and fused ZONOGPT (bottom).

Alg. 1: Structured-zonotope bound propagation

Input: $F = \Phi_T \circ \dots \circ \Phi_1$; input center and radius (x_0, r) ; m shared and q local generators

Output: $(y, \bar{y})$ bounding $F(x)$ for every $\|x - x_0\|_\infty \leq r$

```

1  $Z \leftarrow \mathcal{Z}(x_0, 0, \{\text{diag}(r_s)\}_{s=1}^S, 0)$ ; // Initialize local input generators (§4.1)
2 for  $t \leftarrow 1$  to  $T$  do
3   for  $s \leftarrow 1$  to  $S$  do // Represent  $|e_s| \leq b_s$  with local generators (§4.4)
4      $I_s \leftarrow \{i : (b_s)_i > 0\}$ ;  $L_s \leftarrow [L_s \mid \text{diag}(b_s)_{:, I_s}]$ ;  $b_{s, I_s} \leftarrow 0$ ;
5   if  $\Phi_t$  is an attention block then
6      $Z \leftarrow \text{FusedAttentionTransform}(\Phi_t, Z)$ ; // Bound Attention residual (§4.2)
7   else
8      $Z \leftarrow \text{SequentialMLPTransform}(\Phi_t, Z)$ ; // Bound MLP residual (§4.3)
9    $K \leftarrow \text{SelectShared}(G, m)$ ; // Select at most  $m$  shared generators (§4.4)
10   $b \leftarrow b + \sum_{j \notin K} |G_j|$ ;  $G \leftarrow G_K$ ;
11  for  $s \leftarrow 1$  to  $S$  do // Select at most  $q$  local generators (§4.4)
12     $K_s \leftarrow \text{SelectLocal}(L_s, q)$ ;  $b_s \leftarrow b_s + \sum_{j \notin K_s} |L_{s,j}|$ ;  $L_s \leftarrow L_{s, K_s}$ ;
13 return concretize( $Z$ );
```

Concretizing Eq. 10 gives $[A(X)]_{1,1} \in [-1.24, 1.98]$ (endpoints rounded outward), and cannot verify the property of $[A(X)]_{1,1} > 0$ (Eq. 4).

ZONOGPT instead bounds the attention-residual component as one *fused* block (§4.2):

$$A(x) \in A(c) + J_A(c)G\epsilon + [-R_A, R_A], \quad \|x - x_0\|_\infty \leq r \quad (11)$$

The bottom row of Fig. 2 separates the three contributions in Eq. 11. The output center $A(c)$ computes the attention output at the input center $x = c$. The Jacobian image $J_A(c)G$ carries each input generator through both skip and attention branch, preserving correlations between the two branches. The fused generators encode the nonlinear remainder bounded by R_A . Adding the three terms together gives the final bound, $[A(X)]_{1,1} \in [0.11, 0.69]$, and verifies $[A(X)]_{1,1} > 0$.

4 The ZONOGPT Approach

Alg. 1 describes how ZONOGPT computes bounds for an input interval (x_0, r) through a Transformer F . First, ZONOGPT represents the input interval with local generators (line 1). For each residual component, ZONOGPT represents $|e| \leq b$ with zonotope generators (line 4). It applies fused transformation to the attention block (line 6) and propagates layer-by-layer through the MLP block (line 8). ZONOGPT limits shared generators to m (line 9 to line 10). It also limits local generators to q per token (line 12). Finally, ZONOGPT returns the output bounds (line 13).

The key idea is to preserve dependencies that can tighten later bounds. During bound propagation, ZONOGPT uses *shared generators* for cross-token relations (expensive) and *local generators* for token-specific relations (linear cost in the sequence length), while *intervals* keep the number of generators from growing with network depth.

4.1 Structured Representation

ZONOGPT represents the input using local generators (line 1) because each input coordinate initially affects only its token. Attention later mixes tokens, so ZONOGPT uses shared generators to preserve cross-token relations and an interval to bound the remaining terms independently. This design is inspired by our recent work on preserving cross-dimension relations in input properties [Duong et al., 2026b,c].

Definition 1 (Structured Zonotope). *For S tokens, let c be the center, m shared generators G , q_s local generators L_s at token s , and $b \geq 0$ the interval radius. The represented set is:*

$$\mathcal{Z}(c, G, L, b) = \left\{ c + \sum_{j=1}^m G_j \epsilon_j + \sum_{s=1}^S \text{emb}_s(L_s \epsilon_s^{\text{loc}}) + e : \|\epsilon\|_\infty \leq 1, \|\epsilon_s^{\text{loc}}\|_\infty \leq 1, |e| \leq b \right\} \quad (12)$$

$\text{emb}_s(L_s \epsilon_s^{\text{loc}})$ forms a matrix whose s -th row is $L_s \epsilon_s^{\text{loc}}$ and whose other rows are zero. A generator G_j can change several tokens through the same ϵ_j , while a generator in L_s changes only token s .

The coordinate radius of $Z = \mathcal{Z}(c, G, L, b)$ is $r(Z)$ and interval bounds $[c - r(Z), c + r(Z)]$:

$$r(Z) = \sum_{j=1}^m |G_j| + \sum_{s=1}^S \text{emb}_s \left(\sum_{j=1}^{q_s} |L_{s,j}| \right) + b \quad (13)$$

4.2 Fused Attention Block Transform

For input x , let $A(x)$ denote the attention output with the skip connection. Let n be the LayerNorm output. For head h , let Q_h , K_h , and V_h be the projected queries, keys, and values. Let U_h be the score matrix and P_h the softmax output:

$$\begin{aligned} n &= \text{LN}(x), & Q \| K \| V &= W_{qkv} n + b_{qkv}, & U_h &= Q_h K_h^\top / \sqrt{d_h}, & P_h &= \text{softmax}(U_h) \\ A(x) &= x + W_o \underbrace{[P_1 V_1 \mid \cdots \mid P_H V_H]}_{\text{concat}(P_h V_h)} + b_o \end{aligned} \quad (14)$$

Note that attention contains the related products $QK^\top$ and PV , while softmax couples every key in a row of U_h . Bounding these operations separately loses relations among their terms before adding to the skip branch. ZONOGPT instead propagates each input generator through the complete attention residual and bounds only the remaining nonlinear error.

To construct output bounds for A , ZONOGPT computes the three terms in Eq. 3, e.g., $A(c)$, $H_j + J_B(c)H_j$, and $e_A(c, v)$. First, ZONOGPT computes $A(c)$ simply by evaluating Eq. 14 at $x = c$.

Second, ZONOGPT computes how A changes when the coefficient of H_j varies from zero:

$$x(t) = c + tH_j, \quad \dot{x} = \frac{dx(t)}{dt} \Big|_{t=0} = H_j, \quad \dot{A} = \frac{dA(x(t))}{dt} \Big|_{t=0} = J_A(c)H_j \quad (15)$$

Applying the chain rule along this path gives:

$$\begin{aligned} \dot{n} &= J_{\text{LN}}(c)\dot{x}, & \dot{Q} \| \dot{K} \| \dot{V} &= W_{qkv}\dot{n}, & \dot{U} &= (\dot{Q}K^\top + Q\dot{K}^\top) / \sqrt{d_h}, & \dot{P} &= J_{\text{SM}}(U)\dot{U} \\ \dot{A} &= \dot{x} + W_o \text{concat}(\dot{P}_h V_h + P_h \dot{V}_h) \end{aligned} \quad (16)$$

where each undotted quantity (e.g., Q , K , V , U , and P) denotes its value evaluated at $x = c$. Eq. 15 and Eq. 16 give $J_A(c)H_j = H_j + J_B(c)H_j = \dot{x} + W_o \text{concat}(\dot{P}_h V_h + P_h \dot{V}_h)$.

Third, Taylor expansion identifies error term $e_A(c, v)$ as $A(c + v) - A(c) - J_A(c)v$ (Eq. 2). ZONOGPT then computes R_A satisfying:

$$|e_A(c, v)| = |A(c + v) - A(c) - J_A(c)v| \leq R_A \quad (17)$$

Intuitively, R_A bounds the nonlinear terms introduced by $QK^\top$, softmax, and PV . ZONOGPT computes a bound R_A for the nonlinear attention error e_A using Thm. 1.

Theorem 1 (Attention nonlinear-error bound). *Let A be the attention residual in Eq. 14, and the input zonotope be $Z = \{c + \sum_{j=1}^m H_j \epsilon_j : |\epsilon_j| \leq 1\}$. Let $P = P(c)$ and $V = V(c)$. Assume nonnegative tuples (f_P, R_P, f_V, R_V) denote the bounds on the first-order change (f) and the Taylor remainder (R) of P and V , respectively. satisfy $|J_P(c)(x - c)| \leq f_P$, $|P(x) - P - J_P(c)(x - c)| \leq R_P$, $|J_V(c)(x - c)| \leq f_V$, and $|V(x) - V - J_V(c)(x - c)| \leq R_V$ for every $x \in Z$, then*

$$\underbrace{|A(x) - A(c) - J_A(c)(x - c)|}_{e_A} \leq \underbrace{|W_o| \text{concat}(R_P|V| + PR_V + (f_P + R_P)(f_V + R_V))}_{R_A} \quad (18)$$

Proof sketch. Expanding the attention-value product gives the three terms on the right-hand side of Eq. 18. The triangle inequality bounds these terms by $R_P|V|$, PR_V , and $(f_P + R_P)(f_V + R_V)$, respectively. ZONOGPT derives f_P, R_P, f_V , and R_V in §A.4. $\square$

Thm. 1 bounds the nonlinear attention error with an interval. ZONOGPT improves precision by representing the part caused by the softmax error (e.g., $|W_o| \text{concat}(R_P|V|)$) with a zonotope. Applying the residual Taylor expansion in Eq. 2 to P gives $P(x) = P(c) + J_P(c)(x - c) + e_P$, hence $e_P = P(x) - P(c) - J_P(c)(x - c)$. The entry $(e_P)_{ihj}$ is the softmax error for query token i , attention head h , and key token j . Each row of $P(x)$ and $P(c)$ sums to one (softmax normalizes each row), and each row of $J_P(c)(x - c)$ sums to zero (the derivative of this fixed row sum is zero). Thus, $\sum_j (e_P)_{ihj} = 1 - 1 - 0 = 0$.

Let $\bar{v}_{hj} = W_o^{(h)} V_{hj}(c)$ be the projected center value at key token j . Because the coefficients $(e_P)_{ihj}$ sum to zero, subtracting a_{ih} from every $\bar{v}_{hj}$ does not change their weighted sum. Since $|e_P| \leq R_P$, we write $(e_P)_{ihj} = \xi_{ihj}(R_P)_{ihj}$ with $|\xi_{ihj}| \leq 1$, which gives

$$\begin{aligned} \sum_j (e_P)_{ihj} \bar{v}_{hj} &= \sum_j (e_P)_{ihj} (\bar{v}_{hj} - a_{ih}) \\ &= \sum_j \xi_{ihj} (R_P)_{ihj} (\bar{v}_{hj} - a_{ih}) \end{aligned} \quad (19)$$

Thus, the softmax term is a combination of the generators $(R_P)_{ihj} (\bar{v}_{hj} - a_{ih})$ with coefficients ξ_{ihj} , and their radius $\sum_j (R_P)_{ihj} |\bar{v}_{hj} - a_{ih}|$ depends on a_{ih} . ZONOGPT chooses a_{ih} as an elementwise weighted median of $\bar{v}_{hj}$ with weights $(R_P)_{ihj}$, which minimizes this radius in output.

Finally, ZONOGPT computes a structured zonotope $Z_A = \mathcal{Z}(c_A, G_A, L_A, b_A)$ (Def. 1) for A :

$$\begin{aligned} c_A &= A(c), \quad (G_A)_j = J_A(c)G_j, \\ (L_A)_i &= [(J_A(c) \text{emb}_i(L_{i,k}))_i]_{k=1}^{q_i} \mid [(R_P)_{ihj} (\bar{v}_{hj} - a_{ih})]_{h,j}, \\ (b_A)_i &= [|W_o| \text{concat}(PR_V + (f_P + R_P)(f_V + R_V))]_i + \sum_{s \neq i} \sum_{k=1}^{q_s} |(J_A(c) \text{emb}_s(L_{s,k}))_i| \end{aligned} \quad (20)$$

Thm. 2 establishes soundness of the output zonotope in Eq. 20.

Theorem 2 (Fused attention soundness). *Let the input zonotope be $Z = \mathcal{Z}(c, G, L, 0)$ with local generators $L_i = [L_{i,1}, \dots, L_{i,q_i}]$, and let A be the attention in Eq. 14. Suppose a nonnegative tuple (f_P, R_P, f_V, R_V) satisfies the four bounds in Thm. 1 for every $x \in Z$. For each query token i and head h , choose a_{ih} shared across all key tokens j . Then $Z_A = \mathcal{Z}(c_A, G_A, L_A, b_A)$ defined by Eq. 20 satisfies $A(Z) \subseteq Z_A$.*

Proof sketch. Eq. 19 represents the softmax error with the local generators in L_A . The other nonlinear terms satisfy the bounds in Thm. 1, while b_A bounds the cross-token parts of the local input generator images. The full proof is in §A.5. $\square$

4.3 MLP Block Transform

For input x , the MLP block computes

$$M(x) = x + W_2 g(W_1 \text{LN}(x) + b_1) + b_2 \quad (21)$$

The MLP contains two nonlinear operations with different structures. LayerNorm couples the features within each token through their mean and variance, so ZONOGPT bounds the complete LayerNorm as one fused transform (§4.3.1). GELU acts independently on each coordinate, and the surrounding affine maps are exact (§4.3.2). After LayerNorm, the MLP has no bilinear products such as $QK^\top$ and PV in attention (§4.2). ZONOGPT therefore propagates the affine maps and GELU sequentially before combining the block with the skip connection.

4.3.1 LayerNorm

The mean and variance in LayerNorm depend on the same token features. Bounding these quantities separately loses this relation, so ZONOGPT treats the complete LayerNorm as one fused transform using the Taylor expansion in Eq. 3. At token s , the input zonotope has the form:

$$x_s = c_s + \sum_j (G_j)_s \epsilon_j + L_s \epsilon_s^{\text{loc}} \quad (22)$$

Applying the Taylor expansion to the complete LayerNorm gives

$$\text{LN}(x_s) = \text{LN}(c_s) + \sum_j J_{\text{LN}}(c_s)(G_j)_s \epsilon_j + J_{\text{LN}}(c_s)L_s \epsilon_s^{\text{loc}} + e_{N,s}, \quad |e_{N,s}| \leq (R_N)_s \quad (23)$$

Thus, ZONOGPT computes the LayerNorm output in the form of a structured zonotope as:

$$Z_N = \mathcal{Z}(\text{LN}(c), \{J_{\text{LN}}(c)G_j\}_j, \{J_{\text{LN}}(c)L_s\}_{s=1}^S, R_N) \quad (24)$$

where $c = (c_1, \dots, c_S)$ is the center, and $\text{LN}(c) = (\text{LN}(c_1), \dots, \text{LN}(c_S))$. ZONOGPT derives the remaining term R_N and proves its soundness (Thm. 3) in §A.1.

4.3.2 GELU

The MLP applies tanh-GELU elementwise:

$$g(t) = \frac{t}{2} \left[1 + \tanh \left(\sqrt{\frac{2}{\pi}} (t + 0.044715t^3) \right) \right] \quad (25)$$

For a GELU input zonotope $Z = \mathcal{Z}(c, G, L, b)$, ZONOGPT first computes the coordinate interval $[\ell, u]$. For each coordinate with $\ell_k < u_k$, ZONOGPT separates GELU into the linear term $a_k t$ and the remaining error $g(t) - a_k t$.

$$a_k = \frac{g(u_k) - g(\ell_k)}{u_k - \ell_k}, \quad r_k^{\text{lb}} \leq g(t) - a_k t \leq r_k^{\text{ub}}, \quad t \in [\ell_k, u_k] \quad (26)$$

where the error bounds r^{lb} and r^{ub} are computed in §A.2. With $m = (r^{\text{lb}} + r^{\text{ub}})/2$ and $d = (r^{\text{ub}} - r^{\text{lb}})/2$, Eq. 26 gives the affine GELU bound:

$$g(z) \in a \odot z + m + \{d \odot \epsilon^g : |\epsilon^g| \leq 1\} \quad (27)$$

Applying Eq. 27 to Z gives the GELU output:

$$Z_g = \mathcal{Z}(a \odot c + m, \{a \odot G_j\}_j, \{[\text{diag}(a_s)L_s \mid \text{diag}(d_s)]\}_{s=1}^S, |a| \odot b) \quad (28)$$

GELU bounds are proved sound in Thm. 4 (§A.2).

4.4 Zonotope Conversion and Reduction

In Def. 1, $|e| \leq b$ term means that every coordinate of e varies independently. Before applying block transform, ZONOGPT writes these values as local zonotope generators (line 4). The skip and branch then use the same ϵ for each value while maintaining the original set Z (Lem. 1).

Lemma 1 (Interval-to-zonotope conversion). *A structured zonotope with S tokens satisfies*

$$\mathcal{Z}(c, G, L, b) = \mathcal{Z}\left(c, G, \{[L_s \mid \text{diag}(b_s)]\}_{s=1}^S, 0\right) \quad (29)$$

Proof. The condition $|e_s| \leq b_s$ is equivalent to $e_s = \text{diag}(b_s)\epsilon_s$ for some $|\epsilon_s| \leq 1$. $\square$

After each block transform, ZONOGPT retains at most m shared and q local generators per token (line 9 to line 12). ZONOGPT ranks the generators by summed absolute coefficient magnitude. The selected shared generators remain in G , while $|G_j|$ is added to b for every discarded generator G_j . The resulting zonotope contains the original zonotope (Lem. 2).

Lemma 2 (Generator reduction). *For any set K of retained shared-generator indices,*

$$\mathcal{Z}(c, G, L, b) \subseteq \mathcal{Z}\left(c, G_K, L, b + \sum_{j \notin K} |G_j|\right) \quad (30)$$

Proof. For discarded index j , $|G_j \epsilon_j| \leq |G_j|$ because $|\epsilon_j| \leq 1$. Summing these gives Eq. 30. $\square$

Tab. 1: Benchmark models and instances.

Type	Model	Blocks	Width	Heads	Parameters	Instances
Vision	GPT-2 Small	{2, 4, 8, 12}	768	12	14.2M - 85.1M	720
	GPT-2 Medium	{2, 4, 8, 24}	1024	16	25.3M - 302.4M	720
Text	GPT-2 Small	{2, 4, 8, 12}	768	12	52.8M - 123.7M	720
	GPT-2 Medium	{2, 4, 8, 24}	1024	16	76.7M - 353.8M	720
Total						2880

Tab. 2: Verification results (runtime in seconds). A dash (-) denotes out-of-memory errors.

Dataset	Model	Block	IBP		DEEPZ		$\alpha\beta$ -CROWN		CoVeNN		ZONOGPT-M		ZONOGPT-D		ZONOGPT-F		ZONOGPT-A	
			Solved	Time	Solved	Time	Solved	Time	Solved	Time	Solved	Time	Solved	Time	Solved	Time	Solved	Time
MNIST	GPT-2 Small	2	0	3.6	-	-	-	-	0	5769.8	36	215.3	36	826.2	108	544.5	144	1507.2
		4	0	6.5	-	-	-	-	0	11256.8	0	464.1	0	1423.8	108	2463.5	108	4327.8
		8	0	12.2	-	-	-	-	0	22244.4	0	824.2	0	2942.5	63	5609.2	90	10621.9
		12	0	18.3	-	-	-	-	0	33440.1	0	1232.9	0	4435.2	54	6688.8	72	15862.0
	GPT-2 Medium	2	0	3.6	-	-	-	-	0	10472.1	36	291.9	36	729.3	108	788.3	126	1463.0
		4	0	7.0	-	-	-	-	0	20718.1	0	592.3	0	1506.0	99	2531.2	108	3783.4
		8	0	13.2	-	-	-	-	0	41476.1	0	876.4	0	3048.2	54	4800.8	81	8493.9
		24	0	39.2	-	-	-	-	0	71257.8	0	2612.2	0	9595.7	9	5105.6	9	28909.9
SST	GPT-2 Small	2	0	50.2	-	-	-	-	0	8122.4	24	112.5	50	581.9	87	790.2	94	1962.7
		4	0	77.6	-	-	-	-	0	28785.1	0	343.9	0	1307.2	67	3815.7	81	5260.0
		8	0	61.1	-	-	-	-	0	28194.8	0	649.3	1	2626.8	46	7849.6	61	11036.0
		12	0	68.5	-	-	-	-	0	41694.3	0	1022.4	0	3992.4	72	12677.0	84	17096.6
	GPT-2 Medium	2	0	66.5	-	-	-	-	0	15233.8	42	143.5	66	729.9	91	1027.2	108	2105.2
		4	0	70.3	-	-	-	-	0	27953.7	0	259.7	10	1548.2	74	3018.8	80	4919.5
		8	0	78.8	-	-	-	-	0	53430.5	1	510.1	2	2955.6	56	6984.3	68	10102.5
		24	0	112.2	-	-	-	-	0	149360.1	0	1507.0	0	8472.1	24	18784.8	25	29873.8
Total			0	688.8	-	-	-	-	0	569409.8	139	11657.6	201	46720.8	1120	83479.4	1339	157325.4

4.5 Complexity

Space complexity. Between components, line 9 to line 12 retain at most m shared and q local generators per token for S tokens of width d . The center, b , shared and local generators thus require $O(Sd) + O(mSd) + O(Sqd) = O(Sd(m + q + 1))$. For an attention, Eq. 29 adds $O(Sd^2)$ temporary coefficients before reduction, while the softmax for H heads adds $O(HS^2d)$ and dominates the $O(HS^2)$ attention matrices. The space complexity of ZONOGPT is therefore:

$$O(Sd(m + q + 1) + Sd^2 + HS^2d) \quad (31)$$

Time complexity. Let T_Φ be the time for one full-component Jacobian vector product (JVP), T_{tok} the time for one source-token attention JVP, and T_R the time for computing the nonlinear remainder bound. ZONOGPT has the following time complexity:

$$T_A = O(mT_\Phi + S(q + d)T_{\text{tok}} + T_R), \quad T_{\text{MLP}} = O((m + q + d)T_\Phi + T_R) \quad (32)$$

5 Evaluation

5.1 Experimental Design

Benchmark. We evaluate ZONOGPT on GPT-2-based classifiers that we trained on the MNIST and SST datasets, where SST models use IBP training to improve verifiability [Gowal et al., 2018, Xu et al., 2024b] (§C). We adopted the exact configuration of GPT-2 model [Radford et al., 2019] and utilize its reference implementation from HuggingFace¹. In total, we generated 2880 instances across networks and perturbation radii (Tab. 1). For MNIST, image rows become tokens, and we perturb pixels within an ℓ_∞ ball. For SST, we perturb up to three token embeddings within an ℓ_∞ ball.

Baselines. We compare ZONOGPT with IBP [Wang et al., 2018] (interval), DEEPZ [Singh et al., 2018] (zonotope), $\alpha\beta$ -CROWN [Wang et al., 2021, Zhou et al., 2025] (polytope), and CoVENN [Duong et al., 2025b] (polytope). Among ZONOGPT’s variants, ZONOGPT-D fuses neither blocks (Decomposed), ZONOGPT-A fuses only Attention

¹Hugging Face is a widely adopted open-source library that provides standardized reference implementations for SOTA transformer models. We specifically use their `transformers` package [Wolf et al., 2020].




Fig. 3: Verification and runtime statistics.

(§4.2), ZONOGPT-M fuses only the MLP (§A.3), and ZONOGPT-F fuses both (Full). We exclude DEEPT [Bonaert et al., 2021] and PBVERIFIER [Huang et al., 2026] as they do not support GELU and their post-LayerNorm attention residual (BERT-style) differs from our GPT-based target. All methods run on an NVIDIA GeForce RTX 4090 (24 GB VRAM), an AMD Ryzen Threadripper PRO 5975WX (128 GB RAM).

5.2 RQ1: Overall Verification Performance

Tab. 2 reports verified instances for 2880 instances across MNIST and SST benchmarks. ZONOGPT successfully verifies properties across all network depths, with ZONOGPT-A certifying 46.5% (1339 of 2880) of all instances. In contrast, baseline methods verify none, as they either exceed memory limits (DEEPZ, $\alpha\beta$ -CROWN) or produce loose bounds (IBP, COVENN).

The results validate our block-specific transformations. For the attention block, the fused transform provides the precision necessary to scale to the 24-block GPT-2 Medium. On models with more than two blocks, ZONOGPT-A verifies 40.1% of instances (867 of 2160), whereas unfused-attention variants solve fewer than 1% (ZONOGPT-D solved 13, ZONOGPT-M solved one of 2160). For the MLP block, sequential processing yields higher verified instances than block-level fusion, supporting our analysis in §4.3. Computing the exact affine and elementwise GELU layers sequentially increases verified instances by 19.6% (ZONOGPT-A’s 1339 vs. ZONOGPT-F’s 1120).

The runtimes reflect the computational cost of maintaining precision. Baselines such as IBP and COVENN run quickly, but these loose abstractions fail to verify any instances. In contrast, ZONOGPT incurs more computational overhead to compute tighter bounds, *e.g.*, ZONOGPT-A accrues a substantially higher total runtime (157325s) compared to IBP (689s).

5.3 RQ2: Fused Transformation Effectiveness

To isolate the impact of our fused abstractions, we evaluate performance across varying perturbation radii ($\epsilon \in [10^{-5}, 10^{-3}]$), as shown in Fig. 3a. While all baselines verify zero instances across all tested radii, the ablation of ZONOGPT variants highlights the necessity of the fusion. Variants lacking a fused attention block (ZONOGPT-D and ZONOGPT-M) verify only 146 and 125 instances at the tightest radius ($\epsilon = 10^{-5}$) and drop to near zero by $\epsilon = 10^{-4}$.

Preserving the fused attention allows ZONOGPT-A and ZONOGPT-F to scale to much larger perturbations. Although both perform similarly at $\epsilon = 10^{-5}$, ZONOGPT-A performance is consistently better than ZONOGPT-F across all radii. In particular, at $\epsilon = 10^{-4}$, ZONOGPT-A verifies 305 instances compared to ZONOGPT-F’s 228, and at $\epsilon = 5 \times 10^{-4}$, ZONOGPT-A certifies 68 instances while ZONOGPT-F drops to 9. This highlights the importance of the precision of ZONOGPT-A’s sequential MLP bounds, which are critical as the radius grows.

5.4 RQ3: Runtime Analysis

We evaluate computational efficiency using runtime distributions in Fig. 3b. IBP finishes instantly (under 1s) by relying on cheap but imprecise interval bounds, whereas COVENN incurs massive overhead (median > 150 s) due to computationally expensive CROWN analysis within each block. Nonetheless, both methods fail to maintain the tight bounds required to scale to deep networks.

Within the ZONOGPT variants, runtime predictably scales with abstraction granularity. Fusing the MLP block reduces execution time compared to tracking it sequentially, making ZONOGPT-M (3s) faster than ZONOGPT-D (11s), and ZONOGPT-F (21s) faster than ZONOGPT-A (38s). However, the fused attention abstraction requires slightly more computation to preserve precision, *e.g.*, ZONOGPT-F/ZONOGPT-A run slower than ZONOGPT-M/ZONOGPT-D, respectively. This is because the fused mechanism for attention must capture interactions across multiple paths (including queries, keys, values, and skip connections) into a unified abstract structure.

6 Conclusion

We presented ZONOGPT, an abstract domain for verifying large transformers through structured zonotopes and block-specific fused transformations, and successfully verified standard HuggingFace GPT-2 Medium. Future work includes jointly over-approximating cached attention states for autoregressive decoding [Pope et al., 2023] and adapting to Diffusion [Peebles and Xie, 2023].

References

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.
- Xilie Xu, Keyi Kong, Ning Liu, Lizhen Cui, Di Wang, Jingfeng Zhang, and Mohan Kankanhalli. An llm can fool itself: A prompt-based adversarial attack. In B. Kim, Y. Yue, S. Chaudhuri, K. Fragkiadaki, M. Khan, and Y. Sun, editors, *International Conference on Learning Representations*, volume 2024, pages 2900–2922, 2024a.
- Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.
- Junjie Chu, Yugeng Liu, Ziqing Yang, Xinyue Shen, Michael Backes, and Yang Zhang. Jailbreakradar: Comprehensive assessment of jailbreak attacks against llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21538–21566, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- Yue Huang, Lichao Sun, Haoran Wang, Siyuan Wu, Qihui Zhang, Yuan Li, Chujie Gao, Yixin Huang, Wenhan Lyu, Yixuan Zhang, et al. Position: Trustllm: Trustworthiness in large language models. In *International Conference on Machine Learning*, pages 20166–20270. PMLR, 2024.
- Kun Zhang, Le Wu, Kui Yu, Guangyi Lv, and Dacao Zhang. Evaluating and improving robustness in large language models: a survey and future directions. *arXiv preprint arXiv:2506.11111*, 2025.
- Seongho Joo, Hyukhun Koh, and Kyomin Jung. Harmful prompt laundering: Jailbreaking llms with abductive styles and symbolic encoding. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25500–25535, 2025.
- Robin Jia, Aditi Raghunathan, Kerem Göksel, and Percy Liang. Certified robustness to adversarial word substitutions. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 4129–4142, 2019.
- Kaijie Zhu, Qinlin Zhao, Hao Chen, Jindong Wang, and Xing Xie. Promptbench: A unified library for evaluation of large language models. *Journal of Machine Learning Research*, 25(254):1–22, 2024.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. Bias and fairness in large language models: A survey. *Computational linguistics*, 50(3):1097–1179, 2024.
- Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. Decodingtrust: A comprehensive assessment of trustworthiness in {GPT} models. 2023.

- Zhouxing Shi, Huan Zhang, Kai-Wei Chang, Minlie Huang, and Cho-Jui Hsieh. Robustness verification for transformers. *arXiv preprint arXiv:2002.06622*, 2020.
- Gregory Bonaert, Dimitar I Dimitrov, Maximilian Baader, and Martin Vechev. Fast and precise certification of transformers. In *Proceedings of the 42nd ACM SIGPLAN international conference on programming language design and implementation*, pages 466–481, 2021.
- Pei Huang, Dennis Wei, Omri Isac, Haoze Wu, Min Wu, and Clark Barrett. Parameterized abstract interpretation for transformer verification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 35500–35508, 2026.
- Hai Duong, ThanhVu Nguyen, and Matthew B Dwyer. Generating and checking dnn verification proofs. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025a.
- Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 238–252, 1977.
- Gagandeep Singh, Timon Gehr, Matthew Mirman, Markus Püschel, and Martin Vechev. Fast and effective robustness certification. *Advances in Neural Information Processing Systems*, 31, 2018.
- Shiqi Wang, Huan Zhang, Kaidi Xu, Xue Lin, Suman Jana, Cho-Jui Hsieh, and J. Zico Kolter. Beta-crown: Efficient bound propagation with per-neuron split constraints for neural network robustness verification. In *Advances in Neural Information Processing Systems*, volume 34, pages 29909–29921, 2021.
- Konstantin Kaulen, Tobias Ladner, Stanley Bak, Christopher Brix, Hai Duong, Thomas Flinkow, Taylor T. Johnson, Lukas Koller, Edoardo Manino, ThanhVu H Nguyen, and Haoze Wu. The 6th international verification of neural networks competition (vnn-comp 2025): Summary and results, 2025. URL <https://arxiv.org/abs/2512.19007>.
- Shiqi Wang, Kexin Pei, Justin Whitehouse, Junfeng Yang, and Suman Jana. Formal security analysis of neural networks using symbolic intervals. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1599–1614, 2018. URL <https://dl.acm.org/doi/10.5555/3277203.3277323>.
- Gagandeep Singh, Timon Gehr, Markus Püschel, and Martin Vechev. An abstract domain for certifying neural networks. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–30, 2019.
- Hai Duong, David Shriver, ThanhVu Nguyen, and Matthew B Dwyer. Compositional neural network verification via assume-guarantee reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025b.
- Dennis Wei, Haoze Wu, Min Wu, Pin-Yu Chen, Clark Barrett, and Eitan Farchi. Convex bounds on the softmax function with applications to robustness verification. In *International Conference on Artificial Intelligence and Statistics*, pages 6853–6878. PMLR, 2023.
- Yunruo Zhang, Lujia Shen, Shanqing Guo, and Shouling Ji. Galileo: General linear relaxation framework for tightening robustness certification of transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 21797–21805, 2024.
- Claudio Ferrari, Mark Niklas Mueller, Nikola Jovanović, and Martin Vechev. Complete Verification via Multi-Neuron Relaxation Guided Branch-and-Bound. In *International Conference on Learning Representations*, 2022. doi:10.48550/arXiv.2205.00263.
- Haoze Wu, Omri Isac, Aleksandar Zeljić, Teruhiro Tagomori, Matthew Daggitt, Wen Kokke, Idan Refaeli, Guy Amir, Kyle Julian, Shahaf Bassan, et al. Marabou 2.0: a versatile formal analyzer of neural networks. In *International Conference on Computer Aided Verification*, pages 249–264. Springer, 2024.
- Hai Duong, ThanhVu Nguyen, and Matthew Dwyer. A dpll (t) framework for verifying deep neural networks. *arXiv preprint arXiv:2307.10266*, 2023.
- Hai Duong, Dong Xu, Thanhvu Nguyen, and Matthew B. Dwyer. Harnessing neuron stability to improve dnn verification. *Proc. ACM Softw. Eng.*, 1(FSE), 2024. doi:10.1145/3643765.
- Hai Duong, ThanhVu Nguyen, and Matthew B Dwyer. Neursat: A high-performance verification tool for deep neural networks. In *International Conference on Computer Aided Verification*, pages 409–423. Springer, 2025c.
- Hai Duong and ThanhVu Nguyen. Neursat: Scaling constraint solving for dnn verification (competition contribution). In *International Symposium on AI Verification*, pages 253–259. Springer, 2025.
- Rudy Bunel, P Mudigonda, Ilker Turkaslan, P Torr, Jingyue Lu, and Pushmeet Kohli. Branch and bound for piecewise linear neural network verification. *Journal of Machine Learning Research*, 21(2020), 2020. URL <https://dl.acm.org/doi/10.5555/3455716.3455758>.

- Hai Duong, Le Thanh, and ThanhVu Nguyen. Verifying neural networks with reinforcement learning. *accepted at Advances in Neural Information Processing Systems*, 2026a.
- Duo Zhou, Christopher Brix, Grani A Hanasusanto, and Huan Zhang. Scalable neural network verification with branch-and-bound inferred cutting planes. *arXiv preprint arXiv:2501.00200*, 2024.
- Duo Zhou, Jorge Chavez, Hesun Chen, Grani A. Hanasusanto, and Huan Zhang. Clip-and-verify: Linear constraint-driven domain clipping for accelerating neural network verification. In *Advances in Neural Information Processing Systems*, 2025.
- Hai Duong, Thanh Le, Lam Nguyen, and ThanhVu Nguyen. Verifying Structural Robustness of Deep Neural Network. *Proceedings of the ACM on Software Engineering*, 3(FSE), 2026b.
- Hai Duong, Lam Nguyen, Thanh Le, and ThanhVu Nguyen. Verifying neural network robustness with dual perturbations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27916–27925, 2026c.
- Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Mann, and Pushmeet Kohli. On the effectiveness of interval bound propagation for training verifiably robust models. *arXiv preprint arXiv:1810.12715*, 2018.
- Dong Xu, Nusrat Jahan Mozumder, Hai Duong, and Matthew B Dwyer. Training for verification: Increasing neuron stability to scale dnn verification. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 24–44. Springer, 2024b.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45, 2020.
- Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Jonathan Heek, Kefan Xiao, Shivani Agrawal, and Jeff Dean. Efficiently scaling transformer inference. *Proceedings of machine learning and systems*, 5:606–624, 2023.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4172–4182. IEEE, 2023.
- Huan Zhang, Shiqi Wang, Kaidi Xu, Linyi Li, Bo Li, Suman Jana, Cho-Jui Hsieh, and J Zico Kolter. General cutting planes for bound-propagation-based neural network verification. *Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022. URL <https://dl.acm.org/doi/10.5555/3600270.3600391>.
- Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017. doi:10.1007/978-3-319-63387-9_5.
- Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: a calculus for reasoning about deep neural networks. *Formal Methods in System Design*, 60(1):87–116, 2022. doi:10.1007/s10703-021-00363-7.
- Guy Katz, Derek A Huang, Duligur Ibeling, Kyle Julian, Christopher Lazarus, Rachel Lim, Parth Shah, Shantanu Thakoor, Haoze Wu, Aleksandar Zeljić, et al. The marabou framework for verification and analysis of deep neural networks. In *International Conference on Computer Aided Verification*, pages 443–452. Springer, 2019. doi:10.1007/978-3-030-25540-4_26.
- Stanley Bak. nenum: Verification of ReLU Neural Networks with Optimized Abstraction Refinement. In *NASA Formal Methods Symposium*, pages 19–36. Springer, 2021. doi:10.1007/978-3-030-76384-8_2.
- Hoang-Dung Tran, Diago Manzananas Lopez, Patrick Musau, Xiaodong Yang, Luan Viet Nguyen, Weiming Xiang, and Taylor T Johnson. Star-based reachability analysis of deep neural networks. In *International symposium on formal methods*, pages 670–686. Springer, 2019.
- Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. In *Advances in Neural Information Processing Systems*, volume 31, pages 4939–4948, Montréal, Canada, 2018. Curran Associates, Inc.
- Kaidi Xu, Zhouxing Shi, Huan Zhang, Yihan Wang, Kai-Wei Chang, Minlie Huang, Bhavya Kailkhura, Xue Lin, and Cho-Jui Hsieh. Automatic perturbation analysis for scalable certified robustness and beyond. *Advances in Neural Information Processing Systems*, 33:1129–1141, 2020.

A GPT Transformations

A.1 Fused LayerNorm Transform

Theorem 3 (LayerNorm soundness). *For $\Pi = I - \mathbf{1}\mathbf{1}^\top/d$ and $\tau > 0$, LayerNorm is*

$$\text{LN}(x) = \gamma \odot \frac{\Pi x}{\sqrt{\tau + \|\Pi x\|_2^2/d}} + \beta. \quad (33)$$

For $Z = \mathcal{Z}(c, G, L, 0)$, let H_j range over its shared and embedded local generators. ZONOGPT computes nonnegative vectors r_N and R_N such that every $x \in Z$ satisfies

$$|J_{\text{LN}}(c)(x - c)| \leq r_N, \quad |\text{LN}(x) - \text{LN}(c) - J_{\text{LN}}(c)(x - c)| \leq R_N. \quad (34)$$

The resulting zonotope

$$Z_N = \mathcal{Z}(\text{LN}(c), \{J_{\text{LN}}(c)G_j\}_j, \{J_{\text{LN}}(c)L_s\}_{s=1}^S, R_N) \quad (35)$$

satisfies $\text{LN}(Z) \subseteq Z_N$.

Proof. LayerNorm acts independently on each token, so the proof considers one token. For $x \in Z$, its change from the center and the corresponding centered change are

$$\Delta = x - c = \sum_j H_j \epsilon_j, \quad \bar{\Delta} = \Pi \Delta = \sum_j \underbrace{\Pi H_j}_{\bar{H}_j} \epsilon_j. \quad (36)$$

The centered input and its squared norm at c are

$$z_0 = \Pi c, \quad a_0 = \tau + \|z_0\|_2^2/d. \quad (37)$$

Consequently, LayerNorm at $x = c + \Delta$ is

$$\text{LN}(c + \Delta) = \gamma \odot (z_0 + \bar{\Delta}) \rho(a_0 + \delta a_{\text{lin}} + \delta a_{\text{quad}}) + \beta, \quad (38)$$

where

$$\rho(a) = a^{-1/2}, \quad \delta a_{\text{lin}} = 2z_0^\top \bar{\Delta}/d, \quad \delta a_{\text{quad}} = \|\bar{\Delta}\|_2^2/d. \quad (39)$$

Linear change computation. Differentiating Eq. 38 at $\Delta = 0$ gives

$$J_{\text{LN}}(c)\Delta = \gamma \odot (\bar{\Delta} \rho(a_0) + z_0 \rho'(a_0) \delta a_{\text{lin}}). \quad (40)$$

Because the expression is linear in Δ , substituting Eq. 36 into Eq. 40 gives

$$J_{\text{LN}}(c)\Delta = \sum_j J_{\text{LN}}(c)H_j \epsilon_j. \quad (41)$$

Thus ZONOGPT computes

$$|J_{\text{LN}}(c)\Delta| = |J_{\text{LN}}(c)(x - c)| \leq r_N, \quad r_N = \sum_j |J_{\text{LN}}(c)H_j|, \quad (42)$$

where the inequality follows from $|\epsilon_j| \leq 1$ and the triangle inequality.

Nonlinear error computation. Subtracting $\text{LN}(c)$ and $J_{\text{LN}}(c)\Delta$ from Eq. 38 gives exactly

$$\text{LN}(c + \Delta) - \text{LN}(c) - J_{\text{LN}}(c)\Delta = \gamma \odot (z_0 e_\rho + \bar{\Delta} d_\rho), \quad (43)$$

where

$$\begin{aligned} e_\rho &= \rho(a_0 + \delta a_{\text{lin}} + \delta a_{\text{quad}}) - \rho(a_0) - \rho'(a_0) \delta a_{\text{lin}}, \\ d_\rho &= \rho(a_0 + \delta a_{\text{lin}} + \delta a_{\text{quad}}) - \rho(a_0). \end{aligned} \quad (44)$$

The input generators bound the quantities in Eq. 44 by

$$\begin{aligned} r_{\text{cen}} &= \sum_j |\bar{H}_j|, & |\bar{\Delta}| &\leq r_{\text{cen}}, \\ r_{\text{lin}} &= \sum_j |2z_0^\top \bar{H}_j/d|, & |\delta a_{\text{lin}}| &\leq r_{\text{lin}}, \\ r_{\text{quad}} &= \|r_{\text{cen}}\|_2^2/d, & 0 &\leq \delta a_{\text{quad}} \leq r_{\text{quad}}, \\ r_a &= r_{\text{lin}} + r_{\text{quad}}, & |\delta a_{\text{lin}} + \delta a_{\text{quad}}| &\leq r_a. \end{aligned} \quad (45)$$

Taylor’s theorem also requires a positive lower bound for the argument of ρ . For $\hat{z} = z_0/\|z_0\|_2$ when $z_0 \neq 0$ and $\hat{z} = 0$ otherwise, the generator bounds give

$$\begin{aligned} r_{\parallel} &= \sum_j |\hat{z}^\top \bar{H}_j|, \\ a_{\min} &= \tau + \max \left\{ \max(0, \|z_0\|_2^2/d - r_{\text{lin}}), (\|z_0\|_2 - r_{\parallel})_+^2/d \right\}. \end{aligned} \quad (46)$$

This construction ensures

$$\tau + \|z_0 + t\bar{\Delta}\|_2^2/d \geq a_{\min} \quad \text{for every } t \in [0, 1]. \quad (47)$$

Using $\rho'(a) = -\frac{1}{2}a^{-3/2}$ and $\rho''(a) = \frac{3}{4}a^{-5/2}$, Taylor’s theorem and the mean value theorem yield

$$\begin{aligned} |e_\rho| &\leq \frac{1}{2}a_0^{-3/2}r_{\text{quad}} + \frac{3}{8}a_{\min}^{-5/2}r_a^2, \\ |d_\rho| &\leq \frac{1}{2}a_{\min}^{-3/2}r_a. \end{aligned} \quad (48)$$

Applying these inequalities and $|\bar{\Delta}| \leq r_{\text{cen}}$ to Eq. 43 gives

$$R_N = |\gamma| \odot \left[|z_0| \left(\frac{1}{2}a_0^{-3/2}r_{\text{quad}} + \frac{3}{8}a_{\min}^{-5/2}r_a^2 \right) + \frac{1}{2}r_{\text{cen}}a_{\min}^{-3/2}r_a \right]. \quad (49)$$

Therefore, Eq. 42 gives the first inequality in Eq. 34, and Eq. 49 gives the second. For every $x \in Z = \mathcal{Z}(c, G, L, 0)$, Eq. 34 gives

$$\begin{aligned} \text{LN}(x) &= \text{LN}(c) + \sum_j J_{\text{LN}}(c)G_j\epsilon_j \\ &\quad + \sum_{s=1}^S \text{emb}_s(J_{\text{LN}}(c_s)L_s\epsilon_s^{\text{loc}}) + e_N, \quad |e_N| \leq R_N \end{aligned} \quad (50)$$

Therefore,

$$\text{LN}(Z) \subseteq Z_N = \mathcal{Z}(\text{LN}(c), \{J_{\text{LN}}(c)G_j\}_j, \{J_{\text{LN}}(c_s)L_s\}_{s=1}^S, R_N) \quad (51)$$

The Jacobian images retain every input symbol, while only the nonlinear error e_N enters the interval. $\square$

Remark 1 (Lean mechanization of LayerNorm soundness). *Thm. 3 is fully mechanized in Lean 4 as `layerNorm_image_subset_lnZN` in `ZonoGpt/LayerNorm.lean`, whose only hypothesis is $\tau > 0$. The mechanization follows the proof above. It first shows that the closed-form Jacobian $J_{\text{LN}}(c)$ is the Fréchet derivative of `LayerNorm` (`ZonoGpt/LayerNormDeriv.lean`). For a single token, it then bounds the first-order change by r_N and the Taylor remainder by R_N , using second-order Taylor and mean-value bounds for $a \mapsto a^{-1/2}$ (`ZonoGpt/Rho.lean`, `ZonoGpt/Calculus.lean`); the hypothesis $\tau > 0$ serves only to guarantee $a_{\min} > 0$. Lifting these bounds to all tokens yields the symbol-preserving expansion Eq. 50 (`layerNorm_szero_expansion`), from which the inclusion $\text{LN}(Z) \subseteq Z_N$ follows directly. The proofs proceed step by step, mainly using the tactics `have`, `rw`, `simp`, and `ring`, with `linarith` and `positivity` closing the inequalities. They contain no `sorry` and introduce no axioms.*

A.2 Affine GELU Transform

The MLP requires bounds on tanh-GELU

$$g(t) = \frac{1}{2}t(1 + \tanh h(t)), \quad h(t) = \sqrt{2/\pi}(t + 0.044715t^3) \quad (52)$$

over each preactivation interval $[\ell, u]$. ZONOGPT bounds g on $[\ell, u]$ by a linear lower bound $at + r^{\text{lb}}$ and a linear upper bound $at + r^{\text{ub}}$ that share the coefficient a , where r^{lb} and r^{ub} bound the residual $q(t) = g(t) - at$ (Eq. 59). These bounds hold for every a ; for $\ell < u$, ZONOGPT uses

$$a = \frac{g(u) - g(\ell)}{u - \ell}. \quad (53)$$

The second derivative of g satisfies

$$|g''(t)| \leq \text{sech}^2(h(t)) \left(|h'(t)| + |t| \left(|h'(t)|^2 + \frac{1}{2}|h''(t)| \right) \right), \quad (54)$$

where $h'(t) = \sqrt{2/\pi} (1 + 0.134145t^2)$ and $h''(t) = 0.26829\sqrt{2/\pi}t$. Partition $[\ell, u]$ into $\ell = t_0 < \dots < t_K = u$, and let C_i bound the right-hand side of Eq. 54 on $[t_i, t_{i+1}]$. Linear interpolation of q on this interval gives

$$\begin{aligned} q_i^{\text{lb}} &= \min\{q(t_i), q(t_{i+1})\} - \frac{C_i(t_{i+1} - t_i)^2}{8}, \\ q_i^{\text{ub}} &= \max\{q(t_i), q(t_{i+1})\} + \frac{C_i(t_{i+1} - t_i)^2}{8} \end{aligned} \quad (55)$$

and therefore $q(t) \in [r_{\text{part}}^{\text{lb}}, r_{\text{part}}^{\text{ub}}]$, where

$$r_{\text{part}}^{\text{lb}} = \min_i q_i^{\text{lb}}, \quad r_{\text{part}}^{\text{ub}} = \max_i q_i^{\text{ub}} \quad (56)$$

ZONOGPT also uses the global relation $g(t) = \text{ReLU}(t) - e_{\text{relu}}(t)$, where $e_{\text{relu}}(t) = |t|/(1 + e^{2h(|t|)})$ because h is odd, and

$$0 \leq e_{\text{relu}}(t) \leq e_{\text{max}}. \quad (57)$$

ZONOGPT computes $e_{\text{max}} \approx 0.1703$ numerically. For $|t| \leq 8$, it splits $[0, 8]$ into 8192 equal pieces and bounds e_{relu} on each piece $[t_-, t_+]$ by $t_+/(1 + e^{2h(t_-)})$, since $|t|$ increases and $1/(1 + e^{2h(|t|)})$ decreases in $|t|$. For $|t| \geq 8$, it uses $e_{\text{relu}}(t) \leq 8e^{-2h(8)}$, since $|t|e^{-2h(|t|)}$ decreases there. This computation runs in floating point with a small outward margin rather than directed rounding. A coarser closed-form value follows from $h(|t|) \geq \sqrt{2/\pi}|t|$: $e_{\text{relu}}(t) \leq |t|e^{-2\sqrt{2/\pi}|t|} \leq \sqrt{\pi/2}/(2e) \approx 0.231$. Because $\text{ReLU}(t) - at$ is piecewise affine, its extrema on $[\ell, u]$ occur at ℓ , u , or zero when the interval contains zero. For this set of points E , define

$$r_{\text{global}}^{\text{lb}} = \min_{t \in E} (\text{ReLU}(t) - at) - e_{\text{max}}, \quad r_{\text{global}}^{\text{ub}} = \max_{t \in E} (\text{ReLU}(t) - at) \quad (58)$$

Intersecting the two sound bounds gives

$$r^{\text{lb}} = \max\{r_{\text{part}}^{\text{lb}}, r_{\text{global}}^{\text{lb}}\}, \quad r^{\text{ub}} = \min\{r_{\text{part}}^{\text{ub}}, r_{\text{global}}^{\text{ub}}\} \quad (59)$$

The implementation uses $K = 32$ equal subintervals.

Set $m = (r^{\text{lb}} + r^{\text{ub}})/2$ and $d = (r^{\text{ub}} - r^{\text{lb}})/2$ elementwise. Then every $t \in [\ell, u]$ satisfies

$$g(t) = at + m + d\epsilon^g, \quad |\epsilon^g| \leq 1 \quad (60)$$

For a preactivation zonotope $Z_z = \mathcal{Z}(c_z, G_z, L_z, b_z)$, applying Eq. 60 elementwise produces

$$\begin{aligned} Z_g &= \mathcal{Z}(c_g, G_g, L_g, b_g), & c_g &= a \odot c_z + m, & (G_g)_j &= a \odot (G_z)_j, \\ (L_g)_s &= [a \odot (L_z)_s \mid \text{diag}(d_s)], & (b_g)_s &= |a_s| \odot (b_z)_s \end{aligned} \quad (61)$$

The new columns $\text{diag}(d_s)$ are local because GELU acts independently on every token and coordinate.

Theorem 4 (Affine GELU soundness). *Let $Z_z = \mathcal{Z}(c_z, G_z, L_z, b_z)$ be a GELU input zonotope, and let $[\ell, u]$ contain every coordinate of Z_z . The bounds in Eq. 59 and the output in Eq. 61 satisfy*

$$\{g(z) : z \in Z_z\} \subseteq Z_g. \quad (62)$$

Moreover, the output keeps the input noise symbols: every $z = c_z + \sum_{j=1}^m (G_z)_j \epsilon_j + \sum_{s=1}^S \text{emb}_s((L_z)_s \epsilon_s^{\text{loc}}) + e \in Z_z$ satisfies $g(z) = c_g + \sum_{j=1}^m (G_g)_j \epsilon_j + \sum_{s=1}^S \text{emb}_s((L_g)_s (\epsilon_s^{\text{loc}}, \epsilon_s^g)) + a \odot e$ for some $\|\epsilon^g\|_\infty \leq 1$.

Proof of Thm. 4. On each partition interval, the linear-interpolation error is at most $C_i(t_{i+1} - t_i)^2/8$. The interpolant lies between its endpoint values, which proves Eq. 55. The identity $g = \text{ReLU} - e_{\text{relu}}$ with Eq. 57 proves Eq. 58. Their intersection contains $q(t)$, so centering $[r^{\text{lb}}, r^{\text{ub}}]$ gives Eq. 60. Applying this scalar identity to every coordinate of z gives the symbol-preserving form with the fields in Eq. 61, where $|a \odot e| \leq |a| \odot b_z = b_g$. Therefore, every $g(z)$ belongs to Z_g , which proves Eq. 62. $\square$

Remark 2 (Lean mechanization of affine GELU soundness). *Thm. 4 is mechanized as `gelu_soundness` in `ZonoGpt/GeluSound.lean`. The mechanized bound holds for any $K \geq 1$ partition nodes that cover $[\ell, u]$; the implementation's $K = 32$ equal subintervals and the coefficient a in (53) serve only to tighten the linear lower and upper bounds. The symbol-preserving statement of Thm. 4 is proved as `gelu_szero_symbolic`, which skip connections require because they share noise symbols with earlier generators (`set_inclusion_not_compositional`). The mechanization follows the GELU bound of §A.2 step by step. It first verifies the formulas for g' and g'' , the curvature bound (54), and the closed-form ReLU gap bound $e_{\text{max}} = \sqrt{\pi/2}/(2e)$ (`ZonoGpt/GeluCalc.lean`). On each subinterval, it bounds the residual $q(t) = g(t) - at$ by the linear interpolation error in (55); combining these bounds with the global ReLU bound gives the residual interval $[r^{\text{lb}}, r^{\text{ub}}]$, and hence an affine bound that keeps the noise symbols of the input. The proofs proceed step by step, mainly using the tactics `have`, `rw`, `exact`, and `unfold`, with `linearith` and `positivity` closing the inequalities and `split_ifs` handling the case analysis in the min and max bounds.*

A.3 Fused MLP Transform

For input x , the MLP block computes $M(x) = x + W_2 g(W_1 \text{LN}(x) + b_1) + b_2$. To construct output bounds for M , ZONOGPT computes the three terms of Eq. 3 for the complete MLP residual: $M(c)$, $J_M(c)H_j$, and R_M . First, ZONOGPT computes $M(c)$ by evaluating the MLP residual at $x = c$. Second, ZONOGPT computes $J_M(c)H_j$ along the path $x(t) = c + tH_j$ from Eq. 15:

$$\begin{aligned} n &= \text{LN}(c), & z_c &= W_1 n + b_1, & \dot{n} &= J_{\text{LN}}(c)H_j, & \dot{z} &= W_1 \dot{n} \\ J_M(c)H_j &= H_j + W_2(g'(z_c) \odot \dot{z}) = H_j + J_B(c)H_j \end{aligned} \quad (63)$$

Third, ZONOGPT computes the nonlinear-error bound R_M by Thm. 5.

Theorem 5 (MLP nonlinear-error bound). *Let $M(x) = x + W_2 g(W_1 \text{LN}(x) + b_1) + b_2$ act tokenwise, let $Z = \{c + \sum_j H_j \epsilon_j : |\epsilon_j| \leq 1\}$, and set $z_c = W_1 \text{LN}(c) + b_1$. Suppose a nonnegative tuple (r_N, R_N, C_g) satisfies $|J_{\text{LN}}(c)(x - c)| \leq r_N$ and $|\text{LN}(x) - \text{LN}(c) - J_{\text{LN}}(c)(x - c)| \leq R_N$ for every $x \in Z$. Suppose $|g''(u)| \leq C_g$ elementwise for every u between $z_c - |W_1|(r_N + R_N)$ and $z_c + |W_1|(r_N + R_N)$. Then, for every $x \in Z$,*

$$\underbrace{|M(x) - M(c) - J_M(c)(x - c)|}_{e_M} \leq \underbrace{|W_2 \text{diag}(g'(z_c))W_1|R_N + \frac{1}{2}|W_2| \left(C_g \odot (|W_1|(r_N + R_N))^2 \right)}_{R_M} \quad (64)$$

Proof sketch. Taylor’s theorem splits the omitted error into a transported LayerNorm error and a GELU curvature term. §A.3.1 gives the computation and full proof. $\square$

Finally, ZONOGPT computes the MLP output zonotope $Z_M = \mathcal{Z}(c_M, G_M, L_M, b_M)$ (Def. 1):

$$c_M = M(c), \quad (G_M)_j = J_M(c)G_j, \quad (L_M)_s = [(J_M(c) \text{emb}_s(L_{s,k}))_s]_{k=1}^{q_s}, \quad (b_M)_s = (R_M)_s \quad (65)$$

Thm. 6 establishes soundness of the output zonotope in Eq. 65.

Theorem 6 (Fused MLP soundness). *Let $Z = \mathcal{Z}(c, G, L, 0)$, with $L_s = [L_{s,1}, \dots, L_{s,q_s}]$, and let M be the tokenwise MLP residual above. Suppose a nonnegative tuple (r_N, R_N, C_g) satisfies the hypotheses of Thm. 5 for every $x \in Z$. Then Z_M defined by Eq. 65 satisfies $M(Z) \subseteq Z_M$.*

Proof. The Jacobian images in Eq. 63 retain the input symbols, and the tokenwise MLP keeps each local image at its source token. The nonlinear error satisfies $|e_M| \leq R_M = b_M$ by Thm. 5, so the four fields in Eq. 65 contain every output $M(x)$ for $x \in Z$. $\square$

Remark 3 (Lean mechanization of fused MLP soundness). *Thm. 6 is mechanized as `MlpParams.fused_mlp_sound` in `ZonoGpt/Mlp.lean`, and the variant with LayerNorm pre-composed as `MlpParams.fused_mlp_sound_ln`. The mechanized theorem assumes only the three bounds of Thm. 5 on r_N , R_N , and C_g . The mechanization also proves end-to-end soundness of the sequential MLP, covering the full transform used by the implementation (`MlpParams.Mseq_image_subset_seqMlpOut` in `ZonoGpt/Sequential.lean`). The proof combines two facts. First, the JVP $J_M(c)$ is linear and acts on each token separately, so each input generator maps to exactly one output generator with the same noise symbol, and each local generator stays at its own token. Second, Thm. 5 bounds the remaining error by $R_M = b_M$. Together, they give the symbol-preserving expansion of $M(x)$ (`MlpParams.fused_mlp_symbolic`), from which $M(Z) \subseteq Z_M$ follows. Instantiating r_N and R_N with the LayerNorm radii of Thm. 3, and C_g with $\sup |g''|$, leaves $\tau > 0$ as the only hypothesis (`MlpParams.fused_mlp_sound_ln_sup`). The proofs mainly use the tactics `have`, `rw`, `simp` only, and `exact`, with `linarith` closing the inequalities.*

A.3.1 MLP nonlinear-error computation

Proof of Thm. 5. Fix $x \in Z$ and set $v = x - c$. Set $z_c = W_1 \text{LN}(c) + b_1$. The LayerNorm construction gives the following bounds (Thm. 3):

$$\text{LN}(x) = \text{LN}(c) + J_{\text{LN}}(c)v + e_N, \quad |J_{\text{LN}}(c)v| \leq r_N, \quad |e_N| \leq R_N$$

The change in the GELU preactivation is

$$\begin{aligned} z_{\text{change}} &= W_1(\text{LN}(x) - \text{LN}(c)) \\ &= \underbrace{W_1 J_{\text{LN}}(c)v}_{\text{linear LayerNorm change}} + \underbrace{W_1 e_N}_{\text{LayerNorm nonlinear error}} \end{aligned}$$

Therefore,

$$|z_{\text{change}}| \leq |W_1|r_N + |W_1|R_N = |W_1|(r_N + R_N)$$

For tanh-GELU, set

$$s(u) = \sqrt{2/\pi}(u + 0.044715u^3), \quad g(u) = \frac{1}{2}u(1 + \tanh s(u))$$

The derivatives of the inner function are

$$s'(u) = \sqrt{2/\pi}(1 + 0.134145u^2), \quad s''(u) = 0.26829\sqrt{2/\pi}u$$

Differentiating GELU once gives

$$g'(u) = \frac{1}{2}(1 + \tanh s(u)) + \frac{1}{2}u \operatorname{sech}^2(s(u))s'(u)$$

Differentiating again gives

$$g''(u) = \operatorname{sech}^2(s(u)) \left[s'(u) + \frac{1}{2}us''(u) - u \tanh(s(u))(s'(u))^2 \right]$$

Since $|\tanh(s(u))| \leq 1$,

$$|g''(u)| \leq \operatorname{sech}^2(s(u)) \left(|s'(u)| + |u|(|s'(u)|^2 + \frac{1}{2}|s''(u)|) \right)$$

The preactivation bound places every coordinate of $z_c + z_{\text{change}}$ in

$$[z_c - |W_1|(r_N + R_N), z_c + |W_1|(r_N + R_N)]$$

Bounding the preceding expression over each coordinate interval gives

$$C_g = \sup_{u \in [z_c - |W_1|(r_N + R_N), z_c + |W_1|(r_N + R_N)]} |g''(u)|$$

Apply Taylor's theorem to GELU at z_c :

$$g(z_c + z_{\text{change}}) = g(z_c) + g'(z_c) \odot z_{\text{change}} + e_g$$

The curvature assumption on C_g gives

$$|e_g| \leq \frac{1}{2}C_g \odot |z_{\text{change}}|^2 \leq \frac{1}{2}C_g \odot (|W_1|(r_N + R_N))^2$$

Expanding the complete MLP residual gives

$$\begin{aligned} M(x) - M(c) &= v + W_2(g(z_c + z_{\text{change}}) - g(z_c)) \\ &= v + W_2 \operatorname{diag}(g'(z_c))W_1 J_{\text{LN}}(c)v \\ &\quad + W_2 \operatorname{diag}(g'(z_c))W_1 e_N + W_2 e_g \end{aligned}$$

The first two terms form $J_M(c)v$. Consequently,

$$M(x) - M(c) - J_M(c)v = W_2 \operatorname{diag}(g'(z_c))W_1 e_N + W_2 e_g$$

Taking magnitudes after composing the linear factors yields

$$\begin{aligned} |M(x) - M(c) - J_M(c)v| &\leq |W_2 \operatorname{diag}(g'(z_c))W_1| R_N + |W_2| |e_g| \\ &\leq |W_2 \operatorname{diag}(g'(z_c))W_1| R_N \\ &\quad + \frac{1}{2}|W_2| \left(C_g \odot (|W_1|(r_N + R_N))^2 \right) \end{aligned}$$

The right-hand side is R_M in Eq. 64. $\square$

Remark 4 (Lean mechanization of the MLP nonlinear-error bound). *The computation above is mechanized as `MlpParams.mlp_remainder` in `ZonoGpt/Mlp.lean` using a second-order Taylor bound (`taylor2_bound` in `ZonoGpt/Calculus.lean`). The curvature bound C_g needs to hold only on the interval centered at z_c with radius $|W_1|(r_N + R_N)$. The mechanization defines a computable elementwise curvature bound, `MlpParams.CgSup`, and proves it valid using the global curvature estimate $|g''(u)| \leq C_{\text{sup}}$ from `ZonoGpt/GeluCalc.lean`. The mechanization follows the computation step by step: it writes the preactivation change as W_1 applied to the first-order LayerNorm change plus the LayerNorm remainder, bounds this change elementwise by $|W_1|(r_N + R_N)$, applies the second-order Taylor bound to g on the segment between z_c and the perturbed preactivation, and combines the two error terms into R_M . This Taylor bound requires only that g and g' be differentiable. The bound holds for every τ and for r_N , R_N , and C_g of any sign.*

A.4 Fused Attention Transform

Proof of Thm. 1. Fix $x \in Z$ with $\Delta = x - c$. For each intermediate quantity X , its total change, linear change, and nonlinear error are

$$\delta X = X(x) - X(c), \quad \dot{X} = J_X(c)\Delta, \quad e_X = \delta X - \dot{X} \quad (66)$$

Thus $\dot{X}$ is the linear change retained by the Jacobian, and e_X is the nonlinear error to be bounded.

First, ZONOGPT computes the LayerNorm linear-change and nonlinear-error bounds

$$|\dot{N}| \leq r_N, \quad |e_N| \leq R_N \quad (67)$$

using the construction in §A.1. The affine QKV projection then gives (f_Q, f_K, f_V) from $|W_{qkv}|r_N$ and (R_Q, R_K, R_V) from $|W_{qkv}|R_N$. Hence

$$|\delta Q| \leq T_Q = f_Q + R_Q, \quad |\delta K| \leq T_K = f_K + R_K, \quad |\delta V| \leq T_V = f_V + R_V \quad (68)$$

Next, expand the score matrix $U = QK^\top / \sqrt{d_h}$. Its linear change and nonlinear error are

$$\begin{aligned} \dot{U} &= (\dot{Q}K^\top + Q\dot{K}^\top) / \sqrt{d_h} \\ e_U &= (e_Q K^\top + Q e_K^\top + \delta Q \delta K^\top) / \sqrt{d_h} \end{aligned} \quad (69)$$

The three required bounds follow directly:

$$\begin{aligned} f_U &= (f_Q |K|^\top + |Q| f_K^\top) / \sqrt{d_h} \\ R_U^N &= (R_Q |K|^\top + |Q| R_K^\top) / \sqrt{d_h} \\ R_U^\times &= T_Q T_K^\top / \sqrt{d_h} \end{aligned} \quad (70)$$

Therefore $|\dot{U}| \leq f_U$, $|e_U| \leq R_U^N + R_U^\times$, and $|\delta U| \leq f_U + R_U^N + R_U^\times$.

The score bounds must now pass through softmax. For every nonnegative array r , define

$$[\mathcal{J}_P(r)]_{ij} = P_{ij} \left(r_{ij} + \sum_k P_{ik} r_{ik} \right) \quad (71)$$

so $|J_{\text{SM}}(U)\Delta U| \leq \mathcal{J}_P(|\Delta U|)$. In particular, $f_P = \mathcal{J}_P(f_U)$ bounds $|\dot{P}|$.

To bound the nonlinear softmax error, consider one row $p(t) = \text{softmax}(U + t\delta U)$ and set $M = \|f_U + R_U^N + R_U^\times\|_\infty$. Direct differentiation gives

$$p_j''(t) = p_j(t) [(\delta U_j - \mathbb{E}_{p(t)} \delta U)^2 - \text{Var}_{p(t)}(\delta U)] \quad (72)$$

The bracket has magnitude at most $4M^2$, and $p_j(t) \leq P_j e^{2M}$. Taylor's integral remainder therefore gives

$$|\text{softmax}(U + \delta U) - P - J_{\text{SM}}(U)\delta U| \leq 2P \odot e^{2M} M^2 \quad (73)$$

The retained derivative is $J_{\text{SM}}(U)\dot{U}$. The omitted score error $e_U = \delta U - \dot{U}$ contributes at most $\mathcal{J}_P(R_U^N) + \mathcal{J}_P(R_U^\times)$. Thus the Taylor bound is

$$R_P^T = \mathcal{J}_P(R_U^N) + \mathcal{J}_P(R_U^\times) + 2P \odot e^{2M} M^2 \quad (74)$$

Softmax outputs lie in $[0, 1]$, which independently gives $R_P^I = \max\{P, 1 - P\} + f_P$. Consequently,

$$R_P = \min\{R_P^T, R_P^I\}, \quad T_P = f_P + R_P \quad (75)$$

satisfy $|e_P| \leq R_P$ and $|\delta P| \leq T_P$.

Finally, expand the attention-value product after subtracting its center and linear change:

$$\begin{aligned} e_A &= A(x) - A(c) - J_A(c)\Delta \\ &= W_o \text{concat} \left(\underbrace{e_P V}_{\text{softmax error}} + \underbrace{P e_V}_{\text{value error}} + \underbrace{\delta P \delta V}_{\text{product of both changes}} \right) \end{aligned} \quad (76)$$

The three terms are bounded by

$$|e_P V| \leq R_P |V|, \quad |P e_V| \leq P R_V, \quad |\delta P \delta V| \leq (f_P + R_P)(f_V + R_V) \quad (77)$$

because $P \geq 0$ and $T_V = f_V + R_V$. Concatenating the heads and applying $|W_o|$ gives exactly the bound R_A in Eq. 18. The residual skip and output bias introduce no nonlinear error because both are affine. $\square$

Remark 5 (Lean mechanization of the attention nonlinear-error bound). *Thm. 1 is mechanized as `AttnParams.attention_radius_recipe` in `ZonoGpt/AttentionRecipe.lean`, whose only hypothesis is $\tau > 0$. The quantity M in (73) must be taken per row of the score matrix (`elementwise_M_unsound` in `ZonoGpt/SoftmaxTaylor.lean`). The mechanization follows the computation above step by step: it propagates the LayerNorm radii of Thm. 3 through the projections via $|W|$, through the scores via the exact expansion of $QK^\top$, and through softmax via the radius map $\mathcal{J}_P$ and a row-wise Taylor bound based on the second derivative (72). Each step is an exact identity followed by the triangle inequality; the only analytic inputs are the LayerNorm bounds, which need $\tau > 0$, and the softmax Taylor bound. The computed tuple (f_P, R_P, f_V, R_V) is also proved nonnegative, so it satisfies the nonnegativity condition of Thm. 1. The proofs mainly use the tactics `have`, `rw`, `exact`, `simp` only, and `ring`, with `linarith` and `Finset.sum_le_sum` closing the inequalities.*

A.5 Fused attention soundness

Proof of Thm. 2. Fix $x \in Z$. At the center, $P = P(c)$ and $V = V(c)$. For $x \in Z$,

$$P(x) = P + \delta P, \quad V(x) = V + \delta V$$

The four assumptions give

$$\delta P = J_P(c)(x - c) + e_P, \quad \delta V = J_V(c)(x - c) + e_V, \quad |e_P| \leq R_P, \quad |e_V| \leq R_V \quad (78)$$

Expanding $(P + \delta P)(V + \delta V)$ and subtracting its center and Jacobian terms gives

$$A(x) - A(c) - J_A(c)(x - c) = W_o \text{concat}(e_P V + P e_V + \delta P \delta V) \quad (79)$$

Consider the first term at query token i and head h . Each row of e_P sums to zero, so Eq. 19 gives

$$\sum_j (e_P)_{ihj} \bar{v}_{hj} = \sum_j (e_P)_{ihj} (\bar{v}_{hj} - a_{ih}) \quad (80)$$

For every i, h, j , the bound $|(e_P)_{ihj}| \leq (R_P)_{ihj}$ expresses $(e_P)_{ihj}$ as a coefficient in $[-1, 1]$ multiplied by $(R_P)_{ihj}$. Hence the generators $(R_P)_{ihj}(\bar{v}_{hj} - a_{ih})$ in $(L_A)_i$ contain the complete contribution of $W_o \text{concat}(e_P V)$ at token i .

The remaining two terms satisfy

$$|P e_V + \delta P \delta V| \leq P R_V + (f_P + R_P)(f_V + R_V) \quad (81)$$

because $P \geq 0$, $|\delta P| \leq f_P + R_P$, and $|\delta V| \leq f_V + R_V$. Concatenating the heads and applying W_o bounds their output by

$$|W_o| \text{concat}(P R_V + (f_P + R_P)(f_V + R_V)) \quad (82)$$

which is the nonlinear-error term in b_A .

It remains to place the linear term $J_A(c)(x - c)$ in the output zonotope. The definition of Z gives

$$J_A(c)(x - c) = \sum_j J_A(c) G_j \epsilon_j + \sum_{s=1}^S \sum_{k=1}^{q_s} J_A(c) \text{emb}_s(L_{s,k}) \epsilon_{s,k}^{\text{loc}} \quad (83)$$

The first sum is represented exactly by $(G_A)_j = J_A(c) G_j$ with the original coefficients ϵ_j . For a local input generator at token s , $(L_A)_s$ retains the row at token s . At every token $i \neq s$, the triangle inequality bounds the same generator image by $|(J_A(c) \text{emb}_s(L_{s,k}))_i|$. Summing these bounds over $s \neq i$ and k gives the cross-token term in $(b_A)_i$.

Therefore $c_A = A(c)$ represents the center, G_A and L_A represent the retained terms, and b_A bounds all remaining terms. Thus every $A(x)$ belongs to Z_A , which proves $A(Z) \subseteq Z_A$. $\square$

Remark 6 (Lean mechanization of fused attention soundness). *Thm. 2 is mechanized as `AttnParams.fused_attention_sound_recipe` in `ZonoGpt/AttentionRecipe.lean`, whose only hypothesis is $\tau > 0$. The argument for (19), which builds on the exact expansion `AttnParams.A_sub_sub_dA`, uses the zero-sum property $\sum_j (e_P)_{ihj} = 0$ to replace $\bar{v}_{hj}$ by $\bar{v}_{hj} - a_{ih}$ for any a_{ih} , and writes each $(e_P)_{ihj}$ as $\xi_{ihj}(R_P)_{ihj}$ with $|\xi_{ihj}| \leq 1$. The theorem requires $b = 0$ for the input zonotope, because the linear expansion covers only the generator part of the input (`fused_attention_needs_zero_interval`). The proof separates the algebra from the bounds. It first writes the nonlinear error as the algebraic identity $W_o \text{concat}(e_P V + P e_V + \delta P \delta V)$; the zero-sum property then places the softmax term $e_P V$ into local generators, and b_A bounds the remaining terms. This yields the symbol-preserving form (`AttnParams.fused_attention_symbolic_recipe`), from which $A(Z) \subseteq Z_A$ follows for every choice of a_{ih} ; the weighted median serves only to tighten the bound. The proofs mainly use the tactics `simp` only, `have`, `rw`, `exact`, and `refine`, with `funext` for equalities between sequences.*

B Related Work

DNN verification tools. State-of-the-art DNN verifiers, such as those evaluated in VNN-COMP [Kaulen et al., 2025], combine techniques to improve scalability. $\alpha\beta$ -CROWN [Zhou et al., 2024, 2025, Zhang et al., 2022, Wang et al., 2021] and NEURALSAT [Duong et al., 2024, 2025c] split verification problems and refine bounds on the resulting subproblems. MARABOU [Wu et al., 2024, Katz et al., 2017, 2022, 2019] encodes each NNV instance as a constraint-solving problem and solves with a customized solver, while NNENUM [Bak, 2021] enumerates reachable regions using star sets [Tran et al., 2019]. Note that every NNV tool uses some forms of abstract domains to quickly compute bounds of DNNs given input regions and thus verify problems efficiently.

Abstraction domains. Existing abstractions form a spectrum of speed-accuracy tradeoffs, IBP [Wang et al., 2018] propagates interval bounds layer by layer, making it the fastest method but also the loosest. DEEPZ [Singh et al., 2018] refines this by tracking affine error terms through linear layers, yielding tighter bounds at the cost of higher computational cost. CROWN [Zhang et al., 2018] and DEEPOLY [Singh et al., 2019] apply linear relaxation with full back-substitution to the inputs, producing tighter bounds in a single forward-backward pass. α -CROWN [Xu et al., 2020] extends CROWN by optimizing the slopes of the linear relaxations via gradient ascent, while GCP-CROWN [Zhang et al., 2022] augments α -CROWN with cutting planes on intermediate neurons, pruning infeasible regions, further tightening bounds at the cost of very high computation. However, none of these methods scale to very large and deep networks with millions of parameters, *e.g.*, GPT-2 Medium [Radford et al., 2019].

Transformer verification. Existing methods adapt abstract domains to smaller Transformer models or tighten bounds for individual operations. Shi et al. [2020] verifies small models with simplified LayerNorm. DEEPT extends zonotopes to Transformer operations, but evaluates deeper models with smaller configurations and restricted perturbations [Bonaert et al., 2021]. PBVERIFIER [Huang et al., 2026] and Zhang et al. [2024] improve attention-product bounds, while Wei et al. [2023] improves softmax bounds. COVENN [Duong et al., 2025b] scales through decomposition, whose interval bounds between components lose precision quickly.

Our approach. Prior work remains three limitations: (i) DEEPZ accumulates generators with network size; (ii) existing work bounds softmax and attention products separately, while COVENN’s interval bounds lose relations between components; and (iii) Shi et al. [2020] uses simplified LayerNorm, while DEEPT evaluates smaller configurations with restricted perturbations. ZONOGPT addresses these limitations by (i) reducing shared and local generators after each block to keep space independent of network depth (§4.1, §4.4, and §4.5); (ii) retaining relations between components in its zonotope and bounding the attention residual as one fused transform (§4.1 and §4.2); and (iii) bounding standard LayerNorm and GELU to verify GPT-2 architectures (§4.3.1, §4.3.2, and §5).

C Benchmark Details

Tab. 3: Benchmark configuration.

	MNIST	SST
Task	10-class image classification	Binary sentiment classification
Input sequence	28 projected image rows	GPT-2 tokens, length 32 or 48
Classifier input	Mean-pooled tokens	Final token
Training	Cross-entropy	IBP, from scratch
Perturbation	ℓ_∞ over all pixels	ℓ_∞ over 1 or 3 token embeddings
Perturbation dimension	$28 \times 28 = 784$	768 or 2304 (Small), 1024 or 3072 (Medium)
Radii ε	$10^{-3}, 5 \times 10^{-4}, 10^{-4}, 5 \times 10^{-5}, 10^{-5}$	
Instances	$8 \times 5 \times 4 \times 9 = 1440$	$8 \times 5 \times 18 \times 2 = 1440$

Both tasks use causal GPT-2 blocks at Small (width 768, 12 heads) and Medium (width 1024, 16 heads) sizes. The block counts are 2, 4, 8, and 12 for Small, and 2, 4, 8, and 24 for Medium. For MNIST, a learned projection maps each image row to a token. We select four correctly classified test images per model and radius, clip pixel perturbations to $[0, 1]$, and write one specification per alternative class. For SST, we perturb token embeddings rather than token IDs. We train the SST models from scratch with IBP, since verification-aware training improves the verifiability of the resulting networks [Gowal et al., 2018, Xu et al., 2024b]. We select 18 correctly classified validation sentences per model and radius. For each sentence, we generate two specifications, perturbing either one or three token embeddings. All sampling uses seed 0.

In all experiments, ZONOGPT bounds the combined generator pool by 512 MiB (2^{29} bytes). The shared and local tensors store mSd and qSd float32 values, respectively, and therefore require $4Sd(m+q)$ bytes. The storage limit gives $m+q \leq \lfloor 2^{29}/(4Sd) \rfloor$. The implementation retains at least 64 generators, giving $m+q \leq \max\{64, \lfloor 2^{29}/(4Sd) \rfloor\}$.

Let $C = \max\{64, \lfloor 2^{29}/(4Sd) \rfloor\}$ denote this combined limit, and let $\hat{m}$ and $\hat{q}$ denote the numbers of shared and local generators before each reduction. The values of C are 6,241 and 4,681 for MNIST Small and Medium, and 5,461 and 4,096 for SST Small and Medium, respectively. ZONOGPT computes m and q dynamically during propagation, while keeping $m+q \leq C$ after each reduction. In particular, ZONOGPT retains $m = \min\{\hat{m}, C\}$ shared generators and $q = \min\{\hat{q}, C - m\}$ local generators per token.