Understanding Sensitivity of Differential Attention through the Lens of Adversarial Robustness

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

016

018

019

021

023

024

025

026

027

028

031

033

034

035

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Differential Attention (DA) has been proposed as a refinement to standard attention, suppressing redundant or noisy context through a subtractive structure and thereby reducing contextual hallucination. While this design sharpens task-relevant focus, we show that it also introduces a structural fragility under adversarial perturbations. Our theoretical analysis identifies negative gradient alignment—a configuration encouraged by DA's subtraction—as the key driver of sensitivity amplification, leading to increased gradient norms and elevated local Lipschitz constants. We empirically validate this Fragile Principle through systematic experiments on ViT/DiffViT and evaluations of pretrained CLIP/DiffCLIP, spanning five datasets in total. These results demonstrate higher attack success rates, frequent gradient opposition, and stronger local sensitivity compared to standard attention. Furthermore, depth-dependent experiments reveal a robustness crossover: stacking DA layers attenuates small perturbations via depth-dependent noise cancellation, though this protection fades under larger attack budgets. Overall, our findings uncover a fundamental trade-off: DA improves discriminative focus on clean inputs but increases adversarial vulnerability, underscoring the need to jointly design for selectivity and robustness in future attention mechanisms.

1 Introduction

Recent advances in attention mechanisms have significantly enhanced the representational capacity of deep learning models, driving breakthroughs across vision, language, and multimodal domains. However, these mechanisms can sometimes lead to contextual hallucinations due to the misallocation of attention scores (Huang et al., 2024; Maynez et al., 2020).

The Differential Transformer (Ye et al., 2025), introduced alongside the Differential Attention (DA) mechanism, offers a promising refinement to address such misallocation issues. DA employs two distinct attention maps, A_1 and A_2 , and introduces a subtraction operation, $A_1 - \lambda A_2$, inspired by noise-cancellation techniques in signal processing. This formulation suppresses irrelevant or noisy information while enhancing focus on task-relevant features. The subtraction structure encourages higher attention weights in A_1 for informative regions and lower weights in A_2 for the same regions, yielding sharper and more selective outputs (Figure 1(a)).

Thanks to these advantages, the Differential Transformer effectively mitigates contextual hallucinations in summarization and question answering tasks (Ye et al., 2025). This improvement likely stems from DA's increased selectivity for task-relevant content, aligning with prior findings (Huang et al., 2024) that identify attention misallocation as a key cause of hallucination in standard Transformers.

These properties make DA particularly attractive for safety-critical applications, such as autonomous driving, medical diagnostics, and legal document analysis, where minimizing spurious or incoherent model behavior is essential. These advantages have already inspired a number of follow-up studies (Abid et al., 2025; Hammoud & Ghanem, 2025; Schneider et al., 2025).

At first glance, the subtractive structure of DA may appear inherently beneficial for robustness—by attenuating noisy signals, one might expect it to improve stability against adversarial perturbations as well. However, does this intuition hold? In this work, we rigorously challenge this assumption

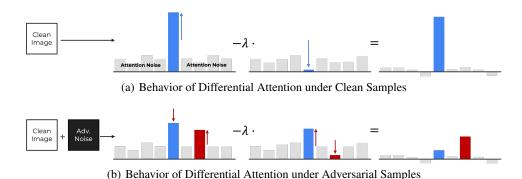


Figure 1: Illustration of the Fragile Principle of Differential Attention. (a) On clean inputs, well-aligned attention maps cancel redundant focus, producing sharp and stable responses. (b) With adversarial perturbations, the gradients (red arrows) of the two attention branches may become negatively aligned, amplifying small input changes and leading to conflicting responses.

and show that the same structure designed to suppress noise can introduce a latent vulnerability. We formalize this phenomenon as the *Fragile Principle* of DA, revealing how subtractive mechanisms amplify input perturbation sensitivity and expose the model to adversarial fragility (Figure 1(b)).

We theoretically show that when the gradients of the two branches are aligned in opposite directions—that is, they point against each other with $\cos\theta < 0$, a condition we call **negative gradient alignment**—the subtraction amplifies the overall gradient norm and increases the local Lipschitz constant. This phenomenon, inherent to DA's subtractive design, not only reduces attention misallocation but also increases sensitivity to adversarial perturbations. Beyond this single-layer fragility, we uncover a **depth-dependent robustness effect**: stacking DA layers enhances robustness to small perturbations via cumulative noise cancellation, though fragility re-emerges at larger budgets.

To validate this, we conduct extensive experiments on DA-based models trained on diverse datasets, evaluating their robustness against adversarial perturbations. The results support our theoretical predictions: DA exhibits higher attack success rates, more frequent negative gradient alignment, and larger local Lipschitz estimates than models with standard attention.

Contributions. Our work makes the following key contributions:

- We present the first theoretical investigation of Differential Attention (DA) from the perspective of adversarial robustness.
- We show that DA's subtraction, while mitigating attention misallocation, structurally induces *negative gradient alignment*, amplifying sensitivity compared to standard attention and revealing a unique adversarial vulnerability (Section 4).
- We develop a depth-dependent analysis, theoretically and empirically demonstrating how cumulative noise cancellation yields partial robustness under small perturbations (Section 4).
- We evaluate DiffCLIP and DiffViT across multiple datasets, showing increased attack success rates, stronger negative gradient alignment, and higher local Lipschitz estimates relative to standard attention (Section 5).

2 RELATED WORK

Attention Robustness and Vulnerability. Transformers (Vaswani et al., 2017) achieve state-of-theart results across domains, yet are more vulnerable to adversarial inputs than CNNs (Bai et al., 2021b; Mo et al., 2022; Jain & Dutta, 2024). ViTs (Dosovitskiy et al., 2021) are especially susceptible to patch attacks due to global receptive fields (Fu et al., 2022; Liu et al., 2023). Transfer-based attacks (Naseer et al., 2021; Ming et al., 2024) and robustness-oriented training (Qin et al., 2022; Wang et al., 2021) show that robustness must be engineered. Differential Attention (DA) suppresses redundant signals through subtraction, but its robustness remains unexplored.

Interpretability and Sparsity in Attention. The explanatory role of attention is contested: some argue it fails to reflect reasoning (Jain & Wallace, 2019; Serrano & Smith, 2019; Bai et al., 2021a), while others promote sparsity or concept alignment (Martins et al., 2020; Pan et al., 2021; Rigotti et al.). Yet sparsity alone often fails to improve attribution (Meister et al., 2021; Pandey et al., 2023). Structured mechanisms such as Structured Attention (Niculae & Blondel, 2017) aim for transparency, but their robustness implications remain unclear.

Lipschitz Behavior and Gradient Sensitivity. Certified robustness often relies on Lipschitz constraints, which bound worst-case changes. Prior work constraints attention layers (Kim et al., 2021; Dasoulas et al., 2021), stabilizes entropy (Zhai et al., 2023), or shows weight decay modulates sensitivity (Kobayashi et al., 2024). Other analyses examine sequence length and normalization (Research, 2023), while LipShiFT (Menon et al., 2025) achieves near-1-Lipschitz ViTs with state-of-the-art certified ℓ_2 robustness. These works suggest Lipschitz properties strongly shape robustness.

Structural Fragility and OOD Sensitivity. Beyond input perturbations, structural vulnerabilities also emerge: CNNs exhibit frequency-specific sensitivity (Tsuzuku et al., 2019; Kanai et al., 2019), topology exposes robustness gaps in attention models (Chen, 2024), and flows assign high likelihood to trivial OODs (Osada et al., 2024). These findings underscore that architecture itself can amplify or dampen robustness. To our knowledge, no prior work studies DA's subtractive design, which we show suppresses redundancy yet paradoxically increases adversarial fragility.

3 PRELIMINARY

In this section, we briefly review the standard attention mechanism and introduce the structure of Differential Attention (DA) (Ye et al., 2025), which forms the basis for our theoretical analysis.

3.1 ATTENTION MECHANISM

Given an input sequence represented as a matrix $X \in \mathbb{R}^{N \times d}$, where N is the sequence length and d is the embedding dimension, the self-attention mechanism computes the output as follows:

$$\operatorname{Attention}(X) = \operatorname{Softmax}\left(\frac{QK^{\top}}{\sqrt{d_k}}\right)V,$$

where the query, key, and value matrices are linear projections: $Q = XW^Q, K = XW^K, V = XW^V$, with $W^Q, W^K, W^V \in \mathbb{R}^{d \times d_k}$ being learned parameters, and d_k denoting the head dimension. The softmax operation normalizes the attention scores across the key dimension, assigning higher weights to tokens that are more relevant to each query.

3.2 DIFFERENTIAL ATTENTION

Differential Attention (DA) (Ye et al., 2025) modifies the standard attention structure by introducing two parallel branches of attention maps, whose outputs are combined in a subtractive manner to suppress redundant or noisy signals. Specifically, DA computes two separate sets of query and key projections:

$$Q_1 = XW_1^Q, \quad K_1 = XW_1^K, \quad Q_2 = XW_2^Q, \quad K_2 = XW_2^K,$$

and a single shared value projection $V = XW^V$ as well as standard attention. Each branch independently computes its own attention map as

$$A_1 = \operatorname{Softmax}\left(\frac{Q_1 K_1^\top}{\sqrt{d_k}}\right), \quad A_2 = \operatorname{Softmax}\left(\frac{Q_2 K_2^\top}{\sqrt{d_k}}\right).$$

The output of DA is then given by:

$$DA(X) = (A_1 - \lambda A_2)V,$$

where λ is a learned or fixed scalar parameter that controls the contribution of the subtractive branch. When λ is trainable, its initial value is set by λ_{init} . The default value of λ_{init} is set to 0.8 through intensive study in (Ye et al., 2025). This structure aims to enhance semantic focus by amplifying meaningful patterns captured by A_1 while suppressing redundant or noisy patterns captured by A_2 .

4 THEORETICAL ANALYSIS: STRUCTURAL SENSITIVITY AND ROBUSTNESS

We now provide a theoretical analysis of the sensitivity of Differential Attention (DA) to input perturbations. While DA is motivated by the goal of suppressing redundant focus through subtraction, we show that the same structure can, under specific conditions, amplify sensitivity and thereby introduce fragility.

We first examine how gradient amplification arises when the two attention branches are negatively aligned—a configuration that DA naturally encourages (Section 4.1). We then extend the analysis to multiple layers, where stacking can yield a depth-dependent cancellation effect (Section 4.2). Finally, we discuss the scope and limitations of these results.

Notation. Let x be a clean input and $x' = x + \xi$ a perturbed input with small perturbation ξ . We analyze the sensitivity of the differential attention map $A_{DA}(x)$ with respect to ξ . We denote by A_{base} the standard (non-differential) attention map.

All proofs of Lemmas and Theorems are provided in Appendix A.

4.1 Fragile Principle: Sensitivity Amplification via the Subtractive Structure

A core feature of DA is its subtractive formulation: $A_{\mathrm{DA}} = A_1 - \lambda A_2$, designed to suppress redundant or noisy attention patterns. For this subtraction to be effective, the two attention maps A_1 and A_2 cannot act independently; they must focus on overlapping regions but with opposing intensities (see Figure 1). In practice, this means that during training the model is implicitly encouraged to assign gradients pointing in opposite directions to A_1 and A_2 within those regions.

Such opposite alignment is not merely a byproduct but a functional necessity: without it, the subtraction would fail to sharpen focus or remove noise. The cancellation mechanism therefore builds in a structural bias toward negative gradient alignment. While this property enables DA to highlight informative regions more clearly than standard attention, it also raises the possibility that the same mechanism could amplify sensitivity to perturbations. Our analysis begins from this structural observation and investigates its consequences for robustness.

Lemma 1. Let θ be the angle between the input gradients of A_1 and A_2 . Then,

$$\|\nabla_{\xi} A_{\text{DA}}\|^2 = \|\nabla_{\xi} A_1\|^2 + \lambda^2 \|\nabla_{\xi} A_2\|^2 - 2\lambda \|\nabla_{\xi} A_1\| \|\nabla_{\xi} A_2\| \cos \theta. \tag{1}$$

The cross-term becomes positive whenever $\cos \theta < 0$, causing gradient amplification.

Theorem 1 (Sensitivity Amplification by Alignment). Let $\rho = \|\nabla_{\mathcal{E}} A_2\| / \|\nabla_{\mathcal{E}} A_1\|$. Then,

$$\|\nabla_{\xi} A_{\text{DA}}\|^2 = \|\nabla_{\xi} A_1\|^2 (1 + \lambda^2 \rho^2 - 2\lambda \rho \cos \theta).$$

Asymptotically,

$$\|\nabla_{\xi} A_{\text{DA}}\| = \begin{cases} (1 - \lambda \rho) \|\nabla_{\xi} A_{1}\| & \text{if } \cos \theta = +1, \\ (1 + \lambda \rho) \|\nabla_{\xi} A_{1}\| & \text{if } \cos \theta = -1. \end{cases}$$
 (2)

Thus, under negative alignment ($\cos \theta < 0$), DA inherently amplifies perturbation sensitivity. This fragility arises not by accident but as a structural byproduct of DA's cancellation goal.

We now compare DA's sensitivity to that of standard attention maps.

Theorem 2 (Relative Sensitivity to Standard Attention). Let $A_{\mathrm{DA}} = A_1 - \lambda A_2$ and A_{base} be standard attention. Define $\gamma = \frac{\|\nabla_{\xi} A_1\|}{\|\nabla_{\xi} A_{\mathrm{base}}\|}$. Then,

$$\frac{\|\nabla_{\xi} A_{\text{DA}}\|}{\|\nabla_{\xi} A_{\text{base}}\|} = \gamma \sqrt{1 + \lambda^2 \rho^2 - 2\lambda \rho \cos \theta},\tag{3}$$

with maximum attained when $\cos \theta = -1$ and minimum when $\cos \theta = +1$.

Theorem 3 (Existence of Amplifying Perturbations). There exists a perturbation ξ such that

$$\frac{\|\nabla_\xi A_{\mathrm{DA}}\|}{\|\nabla_\xi A_{\mathrm{base}}\|} > 1 \quad \text{if and only if} \quad \cos\theta < \frac{1 + \lambda^2 \rho^2 - \gamma^{-2}}{2\lambda\rho}.$$

This condition delineates the (ρ, θ) region where DA exhibits strictly higher sensitivity than standard attention. Since ρ and θ can be adversarially controlled, DA's subtractive design exposes a structural vulnerability.

Implication for Lipschitz constants. To assess how the above gradient amplification affects robustness, we consider the *local Lipschitz constant*, which captures the worst-case sensitivity of the attention map to small perturbations. Relating DA's gradient behavior to Lipschitz continuity reveals that subtraction structurally shifts sensitivity upward, often increasing local Lipschitz values and weakening robustness.

Building on the definition of the local Lipschitz constant,

$$L(x) = \sup_{\xi \neq 0} \frac{\|A(x+\xi) - A(x)\|_2}{\|\xi\|},\tag{4}$$

we derive the following bound:

Lemma 2. Let $L_{DA}(x)$ and $L_{base}(x)$ be the local Lipschitz constants of A_{DA} and A_{base} . Then,

$$\frac{L_{\mathrm{DA}}(x)}{L_{\mathrm{base}}(x)} \leq \gamma \sqrt{1 + \lambda^2 \rho^2 - 2\lambda \rho \cos \theta}.$$

This inequality shows explicitly how the subtraction weight λ and the gradient alignment angle θ together govern DA's local sensitivity. We further analyze the certifiable robustness radius in Appendix B, extending this discussion.

4.2 Depth-Dependent Robustness via Noise Cancellation

Noise cancellation as a structural property of DA. DA is designed with a subtractive form, $A_{\mathrm{DA}} = A_1 - \lambda A_2$, which suppresses components common to both A_1 and A_2 . This noise cancellation effect is independent of gradient alignment: it does not rely on whether ∇A_1 and ∇A_2 are aligned or anti-aligned, but rather on the structural subtraction that systematically reduces shared activations or perturbations. The strength of this effect is governed primarily by the coefficient λ .

Accumulation across depth. When DA layers are stacked, this cancellation compounds: shared noise that survives one layer is more likely to be attenuated again in subsequent layers. In contrast, standard attention lacks such a subtractive mechanism, so its effective perturbation propagation is closer to a simple accumulation across depth. Formally, if $f^{(d)}$ denotes the d-th DA layer with local Lipschitz constant $L_{\mathrm{DA}}^{(d)}$, and $\alpha^{(d)} \in (0,1]$ denotes the average reduction factor from noise cancellation at layer d, then

$$\|\Delta^{(d)}\| \le \alpha^{(d)} L_{\mathrm{DA}}^{(d)} \|\Delta^{(d-1)}\|, \qquad \Delta^{(0)} = \xi.$$

Consequence (upper bound). By recursion, let $F = f^{(D)} \circ \cdots \circ f^{(1)}$ denote the D-layer composition. Then,

$$||F(x+\xi) - F(x)|| \le \left(\prod_{d=1}^{D} \alpha^{(d)} L_{\mathrm{DA}}^{(d)}\right) ||\xi|| = (\bar{\alpha} \, \bar{L}_{\mathrm{DA}})^{D} ||\xi||,$$
 (5)

where $\bar{\alpha}$ and \bar{L}_{DA} denote the geometric means across layers. If $\bar{\alpha}\bar{L}_{DA} < 1$, small perturbations are attenuated with depth, yielding robustness gains that are *specific to DA*.

We note that α is not an independent hyperparameter but rather an emergent factor that is influenced by λ as well as the distribution of activations across layers. Intuitively, larger λ magnifies the subtractive effect, making it statistically more likely that perturbations are partially cancelled across layers, and thus biases α downward. However, the exact value of α is also shaped by training dynamics, and cannot be reduced to a deterministic function of λ .

Theorem 4 (Depth-Dependent Sensitivity of Standard Attention vs. DA). For standard attention, perturbations propagate as

$$\|\Delta^{(D)}\| \le (\bar{L}_{\text{base}})^D \|\xi\|,$$

with no cancellation effect ($\alpha^{(d)} \approx 1$), unlike DA's structural differentiation. For DA, the bound becomes

$$\|\Delta^{(D)}\| \le (\bar{\alpha}\,\bar{L}_{\mathrm{DA}})^D \|\xi\|,$$

where $\bar{\alpha} < 1$ reflects structural noise cancellation.

Corollary 1 (Crossover in Robustness). If $\bar{L}_{DA} > \bar{L}_{base}$ (DA is locally more sensitive) but $\bar{\alpha} < 1$ (nontrivial cancellation), then there exists a depth threshold D^* such that

$$(\bar{\alpha}\,\bar{L}_{\mathrm{DA}})^{D^*} = (\bar{L}_{\mathrm{base}})^{D^*}.$$

For $D < D^*$, DA is more fragile than standard attention; for $D > D^*$, DA becomes asymptotically more robust.

Interpretation. This analysis highlights two independent mechanisms in DA: (i) negative gradient alignment, which locally amplifies fragility, and (ii) noise cancellation, which systematically attenuates shared perturbations and strengthens with depth. The coexistence of these forces explains why DA can appear fragile at shallow depths yet exhibit improved robustness when scaled deeper.

4.3 LIMITATIONS OF THE THEORETICAL ANALYSIS

While our analysis offers a structural understanding of the sensitivity amplification in DA, it is based on a few simplifying assumptions that may not always hold in practice:

Local linear approximation via input gradients. Our derivation relies on local gradient-based sensitivity—i.e., assuming that local behavior of $A_{\rm DA}$ can be well-approximated by its gradient. This is valid under small perturbations, but may not fully capture global nonlinear effects in deep networks.

Layer isolation. We analyze DA in isolation, holding other network layers fixed. In practice, interactions with downstream modules may either mitigate or exacerbate sensitivity.

These assumptions are common in robustness analyses. They clarify how DA's design induces fragility, but also underscore the need for empirical validation (Section 5).

5 EXPERIMENTS

We conduct extensive experiments to validate our theoretical analysis of the Fragile Principle in DA. Specifically, we evaluate: (i) adversarial vulnerability measured by attack success rate (ASR), (ii) the frequency of negative gradient alignment ($\cos(\nabla_\xi A_1, \nabla_\xi A_2) < 0$), and (iii) empirical estimates of the local Lipschitz constant. We also examine how these effects vary with model depth.

5.1 EXPERIMENTAL SETUP

Models. We evaluate two classes of models: (1) ViT and DiffViT trained from for controlled studies, and (2) pre-trained CLIP (ViT-B/16) (Radford et al., 2021) and its DA-based variant DiffCLIP (Hammoud & Ghanem, 2025). The ViT/DiffViT models are lightweight transformers with a D-layer Attention or Differential Attention (DA) module, where $D \in \{1, 2, 4, 8, 12\}$. Unless otherwise stated, we use D=1. The subtraction coefficient λ is initialized at $\lambda_{\text{init}}=0.8$ and updated during training, following (Ye et al., 2025). DiffCLIP extends CLIP by replacing standard attention with DA, while preserving CLIP's contrastive training objective between image and text. This substitution allows DiffCLIP to suppress redundant activations and sharpen focus on discriminative regions. In addition, (Hammoud & Ghanem, 2025) introduce DiffViT, a ViT variant that employs DA.

Datasets. For ViT and DiffViT, we train on CIFAR-10 and CIFAR-100 (Krizhevsky, 2009), consisting of 10 and 100 classes of 32×32 natural images. We also include Tiny ImageNet (Le & Yang, 2015) for ViT and DiffViT with depth 4, with results reported in Appendix. For CLIP and DiffCLIP, we use the MSCOCO 2014 validation set (Lin et al., 2014) (40k images annotated by 80 categories) and the ImageNet-1k validation set (Deng et al., 2009) (50k images across 1,000 categories).

Adversarial Attacks. We consider PGD (Madry et al., 2017), Carlini–Wagner (CW) (Carlini & Wagner, 2017), and AutoAttack (AA) (Croce & Hein, 2020b). PGD is applied both as an adversarial example attack and as an adversarial patch attack. For CLIP and DiffCLIP, perturbations minimize the cosine similarity between the perturbed image encoding $\phi_{\text{visual}}(x+\delta)$ and its text prompt embedding $\phi_{\text{text}}(t)$, where t is a class description (e.g., "a photo of <class>"). For ViT and DiffViT, perturbations instead minimize cross-entropy loss to induce misclassification. PGD and AutoAttack are constrained under ℓ_{∞} norm, while CW is applied under ℓ_{2} . Patch attacks restrict δ to a square of size $w \times w$ at a random image location.

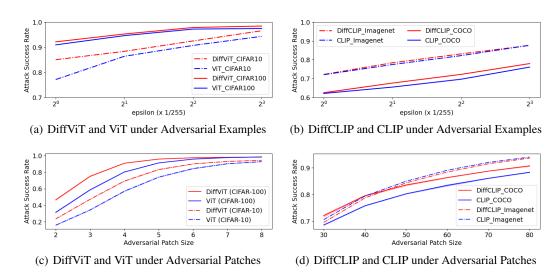


Figure 2: ASR under Adversarial Examples and Patches crafted with PGD. DiffViT and DiffCLIP generally exhibit higher or comparable ASR compared to standard attention, with the gap most pronounced in small-class datasets (CIFAR and COCO), while narrowing on large-scale ones (Imagenet).

Table 1: Attack success rates and clean accuracy of DiffViT on CIFAR-10 when varying λ_{init} . ASR increases with λ_{init} up to 0.8. Beyond that point, ASR drops, suggesting over-cancellation.

$\lambda_{ m init}$	0.5	0.7	0.8 (default)	0.85	0.9	0.95
Accuracy ASR ($\epsilon = 1/255$)	0.8605 0.4074		0.8700 0.8498	0.8567 0.7531		

All experiments are run on a single NVIDIA H100 GPU. Implementation details and hyperparameters are provided in Appendix C.

5.2 ATTACK SUCCESS RATE

To assess adversarial vulnerability, we compute the attack success rate (ASR): $ASR = \frac{\#successful \ attacks}{\#trials}$. Since we focus on untargeted attacks, a perturbation is successful if it changes the prediction relative to the clean input.

Controlled Study on ViT/DiffViT (Single Layer). Figure 2(a) and 2(c) show ASR under adversarial examples and patch attacks, respectively. Across all cases, DiffViT exhibits higher ASR than ViT, consistent with the Fragile Principle. We also vary λ_{init} , which controls the subtraction strength in DA. Table 1 shows ASR on CIFAR-10 increases steadily up to $\lambda_{\text{init}} = 0.8$, after which it declines, suggesting over-subtraction reduces vulnerability.

Controlled Study on ViT/DiffViT (Multiple Layers). Our theory predicts that although a single DA layer embodies the Fragile Principle by amplifying sensitivity, stacking layers reverses this trend through a cancellation effect, producing depth-dependent robustness. Figures 3(c) support this: under CW, deeper DiffViTs require larger perturbations to reach ASR=100% (Fig. 3(c)). Under PGD and AutoAttack, those ASRs are highest at depth 1 but decreases with depth for ϵ =1/255, surpassing ViT's robustness (Fig. 3(a)); at ϵ =4/255, both saturate at high ASR. We report additional experiments on Tiny ImageNet in Appendix D.1, which show the same trend.

Pretrained Models (CLIP/DiffCLIP). Figure 2(b) and 2(d) report ASR for CLIP and DiffCLIP evaluated on COCO and ImageNet. On COCO, DiffCLIP consistently shows higher vulnerability across perturbation budgets and patch sizes, supporting the Fragile Principle. On ImageNet, results are more nuanced: under adversarial examples, DiffCLIP remains slightly more vulnerable, whereas under patch attacks CLIP marginally outperforms DiffCLIP, though the difference is small. This suggests that dataset scale and diversity can modulate the manifestation of DA's fragility.

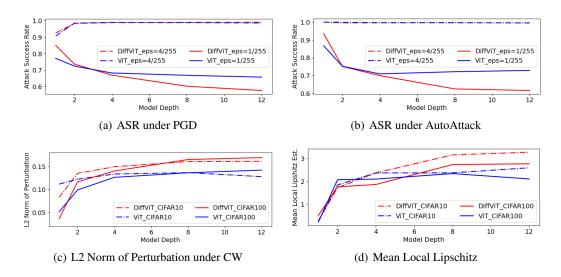


Figure 3: Depth-dependent effects of DA. (a) (b) Under PGD and AutoAttack, DA is fragile at depth 1–2, but ASR drops with depth for ϵ =1/255; at ϵ =4/255 both converge to high ASR. (c) Under CW-L2, deeper models require larger perturbations to reach 100% ASR. (d) Mean local Lipschitz estimates over all layers rise with depth, indicating higher local sensitivity.

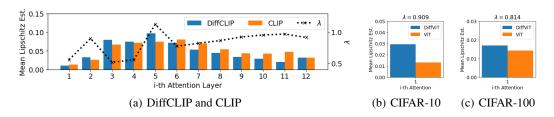


Figure 4: Mean Lipschitz estimates of attention layers under input perturbations. Models incorporating DA exhibit the highest values among all attention layers, particularly at layers with larger λ .

5.3 SENSITIVITY AMPLIFICATION UNDER PERTURBATIONS

We next test whether DA amplifies local sensitivity. We estimate the local Lipschitz constant (4), sampling ξ uniformly from ℓ_{∞} -bounded noise with $\|\xi\| \le 8/255$.

Figure 3(d) reports the *mean local Lipschitz estimate*, computed by averaging the local Lipschitz constants across all layers of each model. We observe that deeper DiffViT models consistently yield higher per-layer Lipschitz values, indicating that local sensitivity increases with depth. Figure 4 further compares ViT/DiffViT and CLIP/DiffCLIP. In all settings, DA-equipped models exhibit significantly higher Lipschitz estimates than baselines, with maxima at layers using larger λ . Interestingly, later layers often display lower sensitivity than earlier ones, suggesting that the cumulative dynamics of DA distribute perturbation effects unevenly across depth rather than simply compounding them. The wider spread of values in DA models highlights their more variable sensitivity across layers.

These results confirm that DA's subtractive design can both amplify and reduce sensitivity, aligning with our theoretical claims. The amplified sensitivity facilitates adversarial perturbations, underscoring risks in safety-critical domains.

5.4 NEGATIVE GRADIENT ALIGNMENT IN DIFFERENTIAL ATTENTION

Finally, we examine whether DA structurally induces negative gradient alignment. For perturbations ξ bounded by $\|\xi\| \le 8/255$, we compute cosine similarity between gradients $\nabla_{\xi} A_1$ and $\nabla_{\xi} A_2$, and record the frequency of $\cos < 0$.

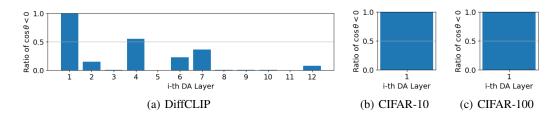


Figure 5: Frequency of negative gradient alignment ($\cos(\nabla A_1, \nabla A_2) < 0$) at each layer.

Figure 5 summarizes results. In DiffCLIP, negative alignment is strongest in the first layer but still present deeper. For CIFAR-trained ViT/DiffViT, even the simplest single-layer variants show dominant negative alignment.

Thus, negative gradient alignment is not a rare anomaly but a structural property of DA, providing direct empirical evidence for the Fragile Principle.

6 DISCUSSION AND LIMITATION

Summary of findings. Our study reveals a structural trade-off at the heart of Differential Attention (DA). While its subtractive design sharpens task-relevant focus and reduces contextual hallucination, it also inherently amplifies sensitivity when the two branches exhibit negative gradient alignment. This *Fragile Principle* manifests as increased local Lipschitz constants, higher attack success rates, and frequent opposing gradient flows. At the same time, our depth-dependent analysis shows that DA models can gain robustness at small perturbation budgets when stacking layers, owing to the cumulative effect of noise cancellation across depth. However, this protection diminishes as perturbation strength grows.

Limitations. Our theoretical analysis relies on local linearization and isolates the attention mechanism from other layers. These assumptions clarify the structural role of subtraction but may not capture full nonlinear interactions. Moreover, our exploration of the subtraction weight λ was limited to initialization; a fuller study of its training dynamics remains open.

Future directions for robustness. Several avenues could mitigate DA's fragility. First, careful tuning of λ offers a trade-off between noise suppression and adversarial robustness, as suggested in Table 1. Second, our results suggest that increasing DA depth itself can serve as a lightweight mitigation: deeper models required larger perturbations under CW and exhibited lower ASR under PGD/AA at small budgets, though this advantage disappears under stronger attacks (Fig. 3). Third, adversarial training with small perturbations reduced ASR (Appendix D.2), showing compatibility with standard defenses. Finally, we view DA as a natural candidate for integration with certified defenses, where its selective subtraction could complement Lipschitz-based guarantees.

Overall, this work highlights both the promise and the fragility of subtractive attention mechanisms, and we hope it inspires future designs that preserve DA's discriminative focus while mitigating its structural vulnerability.

7 CONCLUSION

We analyzed Differential Attention (DA) and revealed a structural fragility that arises under adversarial perturbations. While DA sharpens focus and suppresses contextual hallucination on clean inputs, its subtractive design can amplify local sensitivity via negative gradient alignment. This fragility reflects a fundamental trade-off: DA's discriminative benefits may come at the cost of adversarial vulnerability. We validated this trade-off through gradient alignment analysis, local Lipschitz estimation, and depth-dependent experiments. Importantly, our results show that increasing DA depth can naturally mitigate small perturbations via cumulative noise cancellation, although this effect diminishes under stronger attacks. These insights highlight the need to jointly consider clean performance, adversarial robustness, and architectural depth when designing future attention mechanisms.

Ethics Statement. This work focuses on the theoretical and empirical analysis of Differential Attention (DA) under adversarial perturbations. Our study does not involve human subjects, personal or sensitive data, or any identifiable information. All datasets used in this paper are publicly available benchmarks (CIFAR-10/100, Tiny ImageNet, MSCOCO, ImageNet-1k) that are widely adopted in the research community and collected under ethical and legal guidelines. While our analysis reveals structural vulnerabilities in DA, the goal is to advance understanding of robustness in attention mechanisms and inspire more resilient designs. We do not foresee direct harmful applications of these findings; rather, highlighting DA's fragility is a necessary step toward safer deployment in safety-critical domains such as autonomous driving or medical AI.

Reproducibility Statement. We have taken care to ensure the reproducibility of our results. All model architectures, training setups, and hyperparameters are described in the main text and Appendix C. For theoretical contributions, assumptions are explicitly stated, and proofs of lemmas and theorems are provided in Appendix A. For empirical evaluation, we detail dataset specifications (CIFAR-10/100, Tiny ImageNet, MSCOCO, ImageNet-1k) and attack configurations (PGD, CW, AutoAttack), including perturbation budgets and norm constraints. Results are consistently reported across multiple datasets and model configurations (e.g., depth, λ initialization) to support robustness of conclusions. Figures and tables in the main text and Appendix provide complete experimental outcomes, ensuring transparency and replicability.

REFERENCES

- Mian Muhammad Naeem Abid, Nancy Mehta, Zongwei Wu, and Radu Timofte. Dataformer: Differential additive transformer for lightweight semantic segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 820–831, 2025.
- Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *European conference on computer vision*, pp. 484–501. Springer, 2020.
- Bing Bai, Jian Liang, Guanhua Zhang, Hao Li, Kun Bai, and Fei Wang. Why attentions may not be interpretable? In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pp. 25–34, 2021a.
- Yutong Bai, Jieru Mei, Alan L Yuille, and Cihang Xie. Are transformers more robust than cnns? *Advances in neural information processing systems*, 34:26831–26843, 2021b.
- Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In 2017 ieee symposium on security and privacy (sp), pp. 39–57. Ieee, 2017.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3558–3568, 2021.
- Baiyuan Chen. Robustness evaluation of vision models via topological lipschitz analysis. *arXiv* preprint arXiv:2410.17628, 2024.
- Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *International conference on machine learning*, pp. 2196–2205. PMLR, 2020a.
- Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pp. 2206–2216. PMLR, 2020b.
- George Dasoulas, Kevin Scaman, and Aladin Virmaux. Lipschitz normalization for self-attention layers with application to graph neural networks. In *International Conference on Machine Learning*, pp. 2456–2466. PMLR, 2021.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
 - Yonggan Fu, Shunyao Zhang, Shang Wu, Cheng Wan, and Yingyan Lin. Patch-fool: Are vision transformers always robust against adversarial perturbations? *arXiv preprint arXiv:2203.08392*, 2022.
 - Hasan Abed Al Kader Hammoud and Bernard Ghanem. Diffclip: Differential attention meets clip. *arXiv preprint arXiv:2503.06626*, 2025.
 - Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13418–13427, 2024.
 - Samyak Jain and Tanima Dutta. Towards understanding and improving adversarial robustness of vision transformers. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 24736–24745, 2024. doi: 10.1109/CVPR52733.2024.02336.
 - Sarthak Jain and Byron C Wallace. Attention is not explanation. *arXiv preprint arXiv:1902.10186*, 2019.
 - Shunya Kanai, Kaito Takayama, and Issei Sato. Absum: Learning robustness against adversarial attacks. *arXiv preprint arXiv:1909.08830*, 2019.
 - Hyunjik Kim, George Papamakarios, and Andriy Mnih. The lipschitz constant of self-attention. In *International Conference on Machine Learning*, pp. 5562–5571. PMLR, 2021.
 - Seijin Kobayashi, Yassir Akram, and Johannes Von Oswald. Weight decay induces low-rank attention layers. *Advances in Neural Information Processing Systems*, 37:4481–4510, 2024.
 - Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
 - Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. CS 231N, 7(7):3, 2015.
 - Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision–ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings, part v 13*, pp. 740–755. Springer, 2014.
 - Liang Liu, Yanan Guo, Youtao Zhang, and Jun Yang. Understanding and defending patched-based adversarial attacks for vision transformer. In *Proceedings of the 40th International Conference on Machine Learning*, pp. 21631–21657, 2023.
 - Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
 - André Martins, António Farinhas, Marcos Treviso, Vlad Niculae, Pedro Aguiar, and Mario Figueiredo. Sparse and continuous attention mechanisms. Advances in Neural Information Processing Systems, 33:20989–21001, 2020.
 - Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On faithfulness and factuality in abstractive summarization. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1906–1919, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.173. URL https://aclanthology.org/2020.acl-main.173/.
 - Clara Meister, Stefan Lazov, Isabelle Augenstein, and Ryan Cotterell. Is sparse attention more interpretable? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 296–302, 2021.

- Rohan Menon, Nicola Franco, and Stephan Günnemann. LipshiFT: A certifiably robust shift-based vision transformer. In *ICLR 2025 Workshop: VerifAI: AI Verification in the Wild*, 2025. URL https://openreview.net/forum?id=OfxNKIHfUA.
 - Di Ming, Peng Ren, Yunlong Wang, and Xin Feng. Boosting the transferability of adversarial attack on vision transformer with adaptive token tuning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=sNz7tptCH6.
 - Yichuan Mo, Dongxian Wu, Yifei Wang, Yiwen Guo, and Yisen Wang. When adversarial training meets vision transformers: Recipes from training to architecture. *Advances in Neural Information Processing Systems*, 35:18599–18611, 2022.
 - Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Shahbaz Khan, and Fatih Porikli. On improving adversarial transferability of vision transformers. *arXiv preprint arXiv:2106.04169*, 2021.
 - Vlad Niculae and Mathieu Blondel. A regularized framework for sparse and structured neural attention. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pp. 3338–3348, 2017.
 - Genki Osada, Tsubasa Takahashi, and Takashi Nishide. Understanding likelihood of normalizing flow and image complexity through the lens of out-of-distribution detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 21492–21500, 2024.
 - Bowen Pan, Rameswar Panda, Yifan Jiang, Zhangyang Wang, Rogerio Feris, and Aude Oliva. Iared²: Interpretability-aware redundancy reduction for vision transformers. *Advances in neural information processing systems*, 34:24898–24911, 2021.
 - Shubham Pandey, Ankit Garg, Prateek Jain, and Manik Varma. On the interpretability of attention networks. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pp. 12345–12355, 2023.
 - Yao Qin, Chiyuan Zhang, Ting Chen, Balaji Lakshminarayanan, Alex Beutel, and Xuezhi Wang. Understanding and improving robustness of vision transformers through patch-based negative augmentation. *Advances in Neural Information Processing Systems*, 35:16276–16289, 2022.
 - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.
 - Apple Machine Learning Research. How smooth is attention? *arXiv preprint arXiv:2312.14820*, 2023.
 - Mattia Rigotti, Christoph Miksovic, Ioana Giurgiu, Thomas Gschwind, and Paolo Scotton. Attention-based interpretability with concept transformers. In *International conference on learning representations*.
 - Nadav Schneider, Itamar Zimerman, and Eliya Nachmani. Differential mamba. *arXiv preprint* arXiv:2507.06204, 2025.
 - Sofia Serrano and Noah A Smith. Is attention interpretable? arXiv preprint arXiv:1906.03731, 2019.
 - Yusuke Tsuzuku, Issei Sato, and Masashi Sugiyama. On the structural sensitivity of deep convolutional networks to the directions of fourier basis functions. In *CVPR*, 2019.
 - 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.
 - Boxin Wang, Shuohang Wang, Yu Cheng, Zhe Gan, Ruoxi Jia, Bo Li, and Jingjing Liu. Infobert: Improving robustness of language models from an information theoretic perspective. In 9th International Conference on Learning Representations, ICLR 2021, 2021.

Ross Wightman. Pytorch image models. https://github.com/rwightman/ pytorch-image-models, 2019. Tianzhu Ye, Li Dong, Yuqing Xia, Yutao Sun, Yi Zhu, Gao Huang, and Furu Wei. Differential transformer. In Proceedings of the 13th International Conference on Learning Representations (ICLR 2025), 2025. URL https://openreview.net/forum?id=OvoCm1gGhN. Oral. Shuangfei Zhai, Tatiana Likhomanenko, Etai Littwin, Dan Busbridge, Jason Ramapuram, Yizhe Zhang, Jiatao Gu, and Joshua M Susskind. Stabilizing transformer training by preventing attention entropy collapse. In International Conference on Machine Learning, pp. 40770–40803. PMLR, 2023.

A APPENDIX: PROOF OF THEORETICAL CLAIMS

This section provides complete proofs of the lemmas and theorems presented in Section 4.

A.1 PROOF OF LEMMA 1

Proof. Let θ be the angle between the input gradients of A_1 and A_2 , i.e., $\cos\theta = \frac{\langle \nabla_\xi A_1, \nabla_\xi A_2 \rangle}{\|\nabla_\xi A_1\| \cdot \|\nabla_\xi A_2\|}$. From the definition of DA, $A_{\mathrm{DA}} = A_1 - \lambda A_2$, we compute its gradient with respect to input perturbation ξ : $\nabla_\xi A_{\mathrm{DA}} = \nabla_\xi A_1 - \lambda \nabla_\xi A_2$. Taking the squared norm:

$$\begin{split} \|\nabla_{\xi} A_{\mathrm{DA}}\|^2 &= \|\nabla_{\xi} A_1\|^2 + \lambda^2 \|\nabla_{\xi} A_2\|^2 - 2\lambda \langle \nabla_{\xi} A_1, \nabla_{\xi} A_2 \rangle \\ &= \|\nabla_{\xi} A_1\|^2 + \lambda^2 \|\nabla_{\xi} A_2\|^2 - 2\lambda \|\nabla_{\xi} A_1\| \|\nabla_{\xi} A_2\| \cos \theta. \end{split}$$

A.2 PROOF OF THEOREM 1

Proof. By substituting $\rho = |\nabla_{\xi} A_2|/|\nabla_{\xi} A_1|$ into the result of Lemma 1, we obtain:

$$\|\nabla_{\xi}A_{\mathrm{DA}}\|^2 = \|\nabla_{\xi}A_1\|^2 \left(1 + \lambda^2 \rho^2 - 2\lambda\rho\cos\theta\right).$$

Case 1: If $\cos \theta = +1$ (aligned gradients), then:

$$\|\nabla_{\xi} A_{\mathrm{DA}}\|^2 = \|\nabla_{\xi} A_1\|^2 \left(1 + \lambda^2 \rho^2 - 2\lambda \rho\right) = \|\nabla_{\xi} A_1\|^2 (1 - \lambda \rho)^2.$$

Case 2: If $\cos \theta = -1$ (oppositely aligned gradients), then:

$$\|\nabla_{\xi}A_{\mathrm{DA}}\|^2 = \|\nabla_{\xi}A_1\|^2 \left(1 + \lambda^2 \rho^2 + 2\lambda \rho\right) = \|\nabla_{\xi}A_1\|^2 (1 + \lambda \rho)^2.$$

A.3 PROOF OF THEOREM 2

Proof. Now define $\gamma := \|\nabla_{\xi} A_1\|/\|\nabla_{\xi} A_{\text{base}}\|$. Substituting this into the expression derived in Theorem 1, we get:

$$\|\nabla_{\xi} A_{\mathrm{DA}}\|^2 = \gamma^2 \|\nabla_{\xi} A_{\mathrm{base}}\|^2 (1 + \lambda^2 \rho^2 - 2\lambda \rho \cos \theta).$$

Taking square roots on both sides gives the desired result (3). Finally, note that this expression is minimized when $\cos \theta = +1$ (fully aligned), and maximized when $\cos \theta = -1$ (fully opposed), as the cross term $-2\lambda\rho\cos\theta$ becomes most negative or positive, respectively.

A.4 PROOF OF THEOREM 3

Proof. Let us consider the set of perturbations ξ with $\|\xi\| \le \epsilon$. Since the gradients $\nabla_{\xi}A_1$ and $\nabla_{\xi}A_2$ vary continuously with direction, there exists a direction of ξ within this ϵ -ball such that the angle θ satisfies the required condition. Thus, if the inequality $\cos \theta < \frac{1+\lambda^2\rho^2-\gamma^{-2}}{2\lambda\rho}$ is satisfied for some value of θ (i.e., some ξ), then such a direction ξ with $|\xi| \le \epsilon$ exists, as the set of possible angles θ spans a continuous range over this ball. Since θ varies continuously with ξ , and $\cos \theta$ spans a continuous range, the intermediate value theorem ensures the existence of such a ξ .

A.5 PROOF OF LEMMA 2

Proof. Recall that the local Lipschitz constant of a function f at input x is defined as: $L_f(x) = \sup_{\|\xi\| \neq 0} \frac{\|f(x+\xi)-f(x)\|}{\|\xi\|}$. If f is differentiable at x, then: $L_f(x) \leq \|\nabla_x f(x)\|$. We apply this to the attention maps. Since A_{DA} is differentiable almost everywhere, we have: $L_{\mathrm{DA}}(x) \leq \|\nabla_\xi A_{\mathrm{DA}}(x)\|$. Similarly, $L_{\mathrm{base}}(x) \leq \|\nabla_\xi A_{\mathrm{base}}(x)\|$.

Now, from Theorem 2, we already have the following relationship: $\frac{\|\nabla_{\xi}A_{\mathrm{DA}}(x)\|}{\|\nabla_{\xi}A_{\mathrm{base}}(x)\|} = \gamma\sqrt{1+\lambda^2\rho^2-2\lambda\rho\cos\theta}$. Using the fact that the Lipschitz constant is upper-bounded by the gradient norm: $\frac{L_{\mathrm{DA}}(x)}{L_{\mathrm{base}}(x)} \leq \frac{\|\nabla_{\xi}A_{\mathrm{DA}}(x)\|}{\|\nabla_{\xi}A_{\mathrm{base}}(x)\|} = \gamma\sqrt{1+\lambda^2\rho^2-2\lambda\rho\cos\theta}$.

A.6 PROOF OF THEOREM 4

For standard attention, the layerwise deviation satisfies

$$\|\Delta^{(d)}\| \le L_{\text{base}}^{(d)} \|\Delta^{(d-1)}\|,$$

which unrolls to

$$\|\Delta^{(D)}\| \le \Big(\prod_{d=1}^D L_{\text{base}}^{(d)}\Big) \|\xi\| = (\bar{L}_{\text{base}})^D \|\xi\|.$$

For DA, the subtractive form introduces an additional contraction factor $\alpha^{(d)} \in (0,1]$ per layer, giving

$$\|\Delta^{(d)}\| \le \alpha^{(d)} L_{\mathrm{DA}}^{(d)} \|\Delta^{(d-1)}\|.$$

Thus

$$\|\Delta^{(D)}\| \le \Big(\prod_{d=1}^{D} \alpha^{(d)} L_{\mathrm{DA}}^{(d)}\Big) \|\xi\| = (\bar{\alpha} \, \bar{L}_{\mathrm{DA}})^{D} \|\xi\|,$$

where $\bar{\alpha}$ and \bar{L}_{DA} denote geometric means. Since DA's subtractive structure systematically biases $\alpha^{(d)} < 1$, this establishes the stated bound.

B APPENDIX: ROBUSTNESS CERTIFIABLE RADIUS

Robustness Certifiable Radius. Further, Lipschitz continuity is central to robustness certification. For a classifier f, the certified radius around input x is lower-bounded as:

$$R(x) = \frac{F_y(x) - \max_{i \neq y} F_i(x)}{L_f(x)},$$

where $F_i(x)$ are the class logits, y is the true label, $F_y(x)$ is the logit for the correct class, and $L_f(x)$ is the local Lipschitz constant. The numerator is the classification margin, and the denominator reflects input sensitivity.

The Lipschitz constant of the entire model can be upper-bounded by the product of the per-layer constants $L_f(x) \leq \prod_{i=1}^L L_f^{(i)}(x)$. Since we compare models that differ only in the attention layer, we treat all other $L_f^{(i)}(x)$ as fixed, isolating the contribution of the attention mechanism to the overall sensitivity. Thereby yielding the following theorem about the certified radius.

Theorem 5. Let $R_{\mathrm{DA}}(x)$ and $R_{\mathrm{base}}(x)$ be the certified radius for classifier f incorporating DA and standard attention, respectively. Assume that all per-layer Lipschitz constants $L_f^{(i)}(x)$ except for the attention layer are identical between the two models, isolating the effect of the attention mechanism. Then, the ratio of certified radii satisfies:

$$\frac{R_{\rm DA}(x)}{R_{\rm base}(x)} \ge \frac{\Delta m}{\gamma \sqrt{1 + \lambda^2 \rho^2 - 2\lambda \rho \cos \theta}} \tag{6}$$

where m_{DA} is the classification margin under DA, and $\Delta m = \frac{m_{\mathrm{DA}}}{m_{\mathrm{base}}}$. In particular, under the opposite alignment $\cos\theta = -1$, the bound admits the asymptotic form $\mathcal{O}\left(\frac{\Delta m}{\gamma(1+\lambda\rho)}\right)$.

Proof. Now, using the definition of certified radius and margin ratio $\Delta m := \frac{m_{\mathrm{DA}}}{m_{\mathrm{base}}}$, we compute:

$$\frac{R_{\rm DA}(x)}{R_{\rm base}(x)} = \frac{m_{\rm DA}/L_{\rm DA}(x)}{m_{\rm base}/L_{\rm base}(x)} = \Delta m \cdot \frac{L_{\rm base}(x)}{L_{\rm DA}(x)}. \label{eq:base_equation}$$

Then applying the Lipschitz ratio bound derived in Lemma 2, we finally have the claim (6).

This formulation illustrates a key trade-off in DA: while the subtractive mechanism can sharpen focus and modestly increase classification margins (i.e., $\Delta m > 1$), it also amplifies local sensitivity through the $\gamma \sqrt{1 + \lambda^2 \rho^2 - 2\lambda \rho \cos \theta}$ factor—particularly under negative gradient alignment. Since

this amplification often outweighs the margin gain, the certified robustness deteriorates, as reflected in a shrinking certified radius.

Here, while γ and λ are fixed constants, both Δm , ρ and θ depend on the perturbation direction ξ and can thus be influenced by an adversary. This implies that an attacker may reduce the certified radius by manipulating ρ (e.g., increasing gradient asymmetry) or suppressing the margin $m_{\rm DA}$ through targeted perturbations.

C APPENDIX: IMPLEMENTATION DETAILS

This section describes the implementation details of our experiments to ensure reproducibility.

C.1 TARGET MODELS

DiffCLIP (Hammoud & Ghanem, 2025) extends CLIP (Radford et al., 2021) by replacing standard self-attention with Differential Attention (DA). It retains CLIP's original contrastive training objective between images and text, while using DA to suppress redundant activations and enhance focus on discriminative regions. Additionally, (Hammoud & Ghanem, 2025) introduces DiffViT, a vision transformer variant that incorporates DA into ViT architectures.

DiffCLIP and CLIP. We use the pre-trained weights of DiffCLIP_ViTB16_CC12M released by (Hammoud & Ghanem, 2025), which are trained on CC12M (Changpinyo et al., 2021) datasets. The authors have also publicly released the corresponding CLIP models via their HuggingFace repository. ¹

DiffViT and ViT. For ViT, we adopt the architecture provided in the timm library (Wightman, 2019) and train it on CIFAR-10 and CIFAR-100. We implement DiffViT following the design in (Hammoud & Ghanem, 2025) and train it on the same datasets. We train the models for 100 epochs on CIFAR-10 and 300 epochs on CIFAR-100 using the Adam optimizer with a learning rate of $\eta = 5 \times 10^{-4}$ and a batch size of B = 128.

C.2 ADVERSARIAL ATTACKS

C.2.1 PGD

We evaluate adversarial robustness using both adversarial patch attacks and adversarial examples, optimized via Projected Gradient Descent (PGD) (Madry et al., 2017) under ℓ_{∞} norm constraints.

For CLIP and DiffCLIP, the attack objective is to break the semantic alignment between visual and textual representations. We minimize the cosine similarity between the perturbed visual embedding and its associated text prompt:

$$\min_{\delta} \cos \left(\phi_{\text{visual}}(x+\delta), \phi_{\text{text}}(t) \right),$$

where x is the input image, t is the class prompt (e.g., "a photo of <class>"), and ϕ_{visual} , ϕ_{text} denote the visual and text encoders, respectively.

For ViT and DiffViT on CIFAR-10/100, the attack instead minimizes the (negative) cross-entropy loss to induce misclassification.

In the adversarial patch setting, the perturbation δ is constrained to a fixed-size square patch of dimension $w \times w$ inserted at a random location in x. We set $\epsilon = 1.0$ and vary the patch size w.

C.2.2 AUTOATTACK

AutoAttack is an ensemble adversarial evaluation framework that combines four complementary attacks: APGD-CE, APGD-DLR, FAB (Croce & Hein, 2020a), and Square Attack (Andriushchenko et al., 2020). This combination provides a strong and diverse set of adversarial perturbations for robustness evaluation. We employ the official AutoAttack implementation available at GitHub².

https://huggingface.co/collections/hammh0a/diffclip-67cd8d3b7c6e6ea1cc26cd93

²https://github.com/fra31/auto-attack

C.2.3 CARLINI-WAGNER (CW) ATTACK

The Carlini–Wagner attack (Carlini & Wagner, 2017) is a strong optimization-based adversarial method that searches for minimal perturbations capable of inducing misclassification. Unlike iterative gradient-based methods such as PGD, CW directly formulates the attack as a constrained optimization problem under an ℓ_2 norm. We use a PyTorch implementation based on the open-source repository³.

D APPENDIX: ADDITIONAL EXPERIMENTAL RESULTS

D.1 ASR ON TINYIMAGENET

On Tiny ImageNet (Table 2), we observe a consistent pattern: DiffViT shows higher ASR than ViT across patch sizes and small ℓ_{∞} budgets. The gap is most pronounced at $\epsilon = 1/255$. At larger budgets ($\epsilon = 4/255$), the two models converge, echoing the cancellation–fragility tradeoff seen in CIFAR experiments.

Table 2: ASR on TinyImageNet under PGD Attacks

Model	patch=8	patch=4	patch=2	$\epsilon = 4/255$	$\epsilon = 2/255$	$\epsilon = 1/255$	$\epsilon = 0.5/255$
DiffViT ViT	0.6465 0.6139	0.4556 0.4072	0.1834 0.1498	0.6318 0.6490	0.3866 0.3857	0.2822 0.2309	0.2364 0.1866

D.2 ADVERSARIAL TRAINING

In addition to balancing the structural trade-off inherent to DA, complementary robustness strategies can be applied. Adversarial training offers a complementary strategy that orthogonal to DA's subtractive structure.

Table 3 report the result. The adversarial training with small perturbations successfully reduced ASR without harming classification accuracy. However, beyond $\epsilon=0.5/255$, it drastically drops the clean accuracy. Therefore, adversarial training is a candidate technique to increase the adversary robustness, but is not a silver-bullet.

Table 3: Attack success rates and clean accuracy of vanilla DiffViT and DiffViT with adversarial training when varying ϵ on CIFAR-10.

	Accuracy	ASR (ϵ =0.25/255)	ASR (ϵ =0.5/255)	ASR (ϵ =1/255)
DiffViT (CIFAR-10)	0.8700	0.3310	0.6367	0.8498
+ Adv. Train (ϵ =0.25/255)	0.9518	0.2537	0.6434	0.8824
+ Adv. Train (ϵ =0.5/255)	0.8931	0.1787	0.4440	0.7676
+ Adv. Train (ϵ =1/255)	0.7384	0.0982	0.2089	0.4302

E LLM USAGE

We used large language models (LLMs) as writing assistants to improve clarity and readability. They were not involved in generating ideas, conducting experiments, or producing results. All technical content and contributions are solely by the authors.

³https://github.com/kkew3/pytorch-cw2