

# 000 001 002 003 004 005 006 007 008 009 010 THEORETICALLY UNDERSTANDING THE HIDDEN AD- VERSARIAL PRICE OF LOW-RANK ADAPTATION

005 **Anonymous authors**

006 Paper under double-blind review

## ABSTRACT

011 Low-rank adaptation (LoRA) has emerged as a prominent parameter-efficient fine-  
012 tuning (PEFT) method for large pre-trained models, enabling strong downstream  
013 performance with minimal parameter updates. While LoRA is known to outper-  
014 form head-only fine-tuning in terms of clean accuracy, its impact on adversarial  
015 robustness remains largely unexplored. In this work, and to the best of our knowl-  
016 edge, we present the first theoretical analysis of LoRA’s adversarial robustness,  
017 comparing it to that of head-only fine-tuning. We formalize the notion of ex-  
018 pected adversarial robustness and derive upper bounds demonstrating that, despite  
019 its superior clean performance, LoRA can be inherently less robust than head-only  
020 tuning due to the additional degrees of freedom introduced by its low-rank com-  
021 ponents. We further study the influence of LoRA’s initialization scheme and show  
022 that simple changes in the initialization distribution of the low-rank matrix can  
023 significantly affect robustness. Finally, we support our theoretical findings with  
024 extensive experiments on both vision and language benchmarks under standard  
025 adversarial attacks. Our results provide a principled understanding of the trade-  
026 offs between parameter efficiency, clean performance, and adversarial robustness  
027 in commonly used fine-tuning strategies.

## 028 1 INTRODUCTION

030 Deep learning has led to significant breakthroughs across  
031 multiple domains, notably in computer vision (Dosovitskiy et al., 2021; Liu et al., 2021) and natural language  
032 processing (Devlin et al., 2019; Radford et al., 2019; Jiang  
033 et al., 2023), where foundation models have become cen-  
034 tral. These models, typically based on Transformer archi-  
035 tectures (Vaswani et al., 2017), are pre-trained on large-  
036 scale datasets using auxiliary self-supervised tasks, en-  
037 abling them to learn transferable representations. When  
038 fine-tuned, they achieve state-of-the-art performance on a  
039 wide range of downstream tasks. However, these foun-  
040 dation models are often extremely large, encompassing a  
041 lot of parameters that could range from millions to bil-  
042 lions, which makes full fine-tuning both computationally  
043 expensive and impractical for many users. As a result,  
044 parameter-efficient fine-tuning (PEFT) (Han et al., 2024)  
045 strategies have gained attention. A common approach is to freeze the pretrained model and only  
046 train a lightweight classification or regression head (Kornblith et al., 2019; Chen et al., 2020). While  
047 efficient, this method often yields suboptimal downstream performance. To address this, Low-Rank  
048 Adaptation (LoRA) (Hu et al., 2022) has emerged as a leading PEFT technique. LoRA introduces  
049 learnable low-rank matrices into the model’s weight structure, allowing it to adapt to downstream  
050 tasks while updating only a small subset of parameters. Empirically, LoRA often closely approaches  
051 the performance of full fine-tuning, making it a practical alternative for resource-constrained envi-  
052 ronments.

053 In parallel to advancements in finetuning methods, adversarial robustness remains a pressing chal-  
lenge in deep learning. Neural networks are known to be vulnerable to small, carefully crafted




Figure 1: Clean/attacked accuracy on ImageNet subject to FGSM for a ViT.

054 perturbations that can cause severe misclassifications, even when these perturbations are imper-  
 055 ceptible to humans (Goodfellow et al., 2015). This vulnerability raises concerns in safety-critical  
 056 applications such as autonomous vehicles, healthcare, and finance. While extensive research has  
 057 been conducted on adversarial attack strategies (Tramer et al., 2020; Costa et al., 2024; Biggio et al.,  
 058 2013) and defense mechanisms (Madry et al., 2017; Akhtar et al., 2021), the relationship between  
 059 finetuning strategies and adversarial robustness remains underexplored. In particular, the majority  
 060 of theoretical work has studied how LoRA’s performance is influenced by hyperparameters such  
 061 as rank (Kalajdzievski, 2023), learning rate (Hayou et al., 2024b), and initialization (Hayou et al.,  
 062 2024a), while no theoretical work to date and to our knowledge has rigorously analyzed the impact  
 063 of LoRA on adversarial robustness. Preliminary empirical evidence (such as observed in Figure 1)  
 064 suggests LoRA may influence robustness, but a formal understanding of this phenomenon is lacking.  
 065

066 In this work, we aim to bridge the gap by investigating how LoRA-based fine-tuning affects ad-  
 067 versarial robustness, specifically in comparison to fine-tuning using only a classification or regres-  
 068 sion head. While it is well established that LoRA outperforms head-only tuning in terms of clean  
 069 accuracy, it remains unclear whether this gain comes at the cost of reduced robustness under ad-  
 070 versarial attacks. To address this, we begin by formalizing the notion of expected adversarial ro-  
 071 bustness, which we then use to theoretically analyze and compare the robustness of head-only and  
 072 LoRA-based fine-tuning. Our analysis leads to an upper bound suggesting that head-only fine-tuning  
 073 exhibits stronger adversarial robustness than LoRA, primarily due to the additional parameters in-  
 074 troduced by the low-rank adaptation layers. To further understand the influence of LoRA’s design  
 075 choices, we examine how its initialization scheme impacts robustness. In standard LoRA training,  
 076 one of the low-rank matrices is initialized randomly while the other is set to zero; prior work (Hayou  
 077 et al., 2024a) has shown that initializing  $B$  to zero and  $A$  randomly typically yields better clean ac-  
 078 curacy. We extend this line of investigation by studying how varying the initialization of  $A$  affects  
 079 adversarial robustness and demonstrates that such a simple change can narrow the robustness gap  
 080 between LoRA and head-only tuning. Finally, we empirically validate our theoretical findings on  
 081 both vision and language benchmarks using standard adversarial attacks across multiple datasets.  
 082 Our overall contributions can be summarized as follows:  
 083

- Using a formal notion of expected adversarial robustness, we theoretically show that head-  
 084 only fine-tuning offers higher expected adversarial robustness than LoRA, due to the addi-  
 085 tional degrees of freedom introduced by LoRA’s low-rank matrices.
- We analyze how LoRA’s initialization scheme, particularly the initialization of its low-rank  
 086 matrix  $A$ , and we consequently provide new additional insights on the choice of initial  
 087 distribution, which could reduce the robustness gap between LoRA and head-only fine-  
 088 tuning.
- We validate our theoretical findings through extensive experiments on vision and language  
 089 tasks, using standard adversarial attacks and multiple benchmark datasets.

## 2 RELATED WORK

092 **Parameter-Efficient Fine-Tuning.** Most pre-trained models today are based on the Transformer  
 093 architecture (Vaswani et al., 2017). Fully fine-tuning these large models for downstream tasks is of-  
 094 ten computationally expensive due to the sheer number of parameters, resulting in high memory and  
 095 compute requirements. Parameter-Efficient Fine-Tuning (PEFT) aims to address this challenge by  
 096 introducing a small number of trainable parameters, enabling efficient adaptation without updating  
 097 the entire model. A simple approach is to fine-tune only the task-specific head, which reduces re-  
 098 source usage but often degrades performance. As an alternative, Low-Rank Adaptation (LoRA) (Hu  
 099 et al., 2022), and its variants (Dettmers et al., 2023; Kopiczko et al., 2024; Hayou et al., 2024b; Li  
 100 et al., 2024), inject a small set of trainable parameters into each layer of the frozen Transformer  
 101 backbone, offering a better trade-off between parameter efficiency and downstream performance.  
 102

103 **Initialization of LoRA.** The initialization of the low-rank matrices in LoRA has recently received  
 104 increased attention. Since the product of the two matrices is typically initialized to zero to preserve  
 105 the behavior of the pre-trained model at the start of fine-tuning, various strategies have been pro-  
 106 posed for initializing the non-zero matrix. Recent analysis (Hayou et al., 2024a) shows that this  
 107 choice significantly influences optimization, with initializing  $B$  to zero and  $A$  randomly yielding

108 better average performance. AMT (Yang et al., 2024a) proposes an SVD-based initialization, aligning  
 109 LoRA adapters with principal subspaces of the original weights to improve robustness under  
 110 adversarial tuning. DoRA (Liu et al., 2024) further decomposes pre-trained weights into magnitude  
 111 and direction, restricting LoRA updates to the directional component, leading to improved perfor-  
 112 mance and stability compared to standard LoRA.

113 **Adversarial Robustness and LoRA.** Most prior work on LoRA has focused on its effectiveness for  
 114 downstream task performance. However, recent studies have begun to explore the relationship be-  
 115 tween fine-tuning and adversarial robustness. In particular, works such as (Turbal et al., 2024; Yang  
 116 et al., 2024b) empirically investigate the robustness of large language models in transfer settings.  
 117 In the same direction, AutoLoRA (Xu et al., 2024) and ADV-LoRA (Wu et al., 2025) incorporate  
 118 adversarial training, one of the most established techniques in robustness research, into the LoRA  
 119 framework to enhance resilience. Despite these empirical advances, a theoretical understanding of  
 120 how LoRA and its associated hyperparameters affect adversarial robustness remains lacking. This  
 121 work aims to bridge that gap by developing a general theoretical framework linking LoRA to adver-  
 122 sarial robustness, offering both theoretical and empirical insights that improve model resilience and  
 123 open new research perspective.

### 3 PRELIMINARIES

127 In this section, we start by introducing some fundamental concepts that will be used afterwards in  
 128 our work. Afterward, we formulate our problem setup, which will be considered in our analysis.

129 **Transformer-based Models.** Let  $X \in \mathcal{X} \subseteq \mathbb{R}^{n \times d}$  denote a sequence of  $n$  tokens, where each  
 130 token  $x_i \in \mathbb{R}^d$ . The backbone of a Transformer  $h : \mathcal{X} \subseteq \mathbb{R}^{n \times d} \rightarrow \mathcal{Z} \subseteq \mathbb{R}^{n \times d}$ , as introduced  
 131 in (Vaswani et al., 2017), is the *self-attention* mechanism, which computes a weighted combination  
 132 of all token representations. Specifically, given learnable *query*, *key*, and *value* parameter matrices  
 133  $W^Q, W^K, W^V \in \mathbb{R}^{d \times (d/H)}$ , the output of a single *attention head* AH for input  $X$  is defined as:

$$134 \quad \text{AH}(X) = \text{softmax} \left( \frac{(XW^Q)(XW^K)^\top}{\sqrt{d/H}} \right) (XW^V), \quad (1)$$

135 where  $H$  denotes the number of parallel attention heads and  $d/H$  is the dimension per head. In  
 136 practice, multiple attention heads  $\text{AH}_i$  are computed in parallel, then concatenated and projected  
 137 using a learnable weight matrix  $W^O \in \mathbb{R}^{d \times d}$ , yielding the multi-head attention (MHA) operation:

$$138 \quad \text{MH}(X) = \text{concat}(\text{AH}_1(X), \text{AH}_2(X), \dots, \text{AH}_H(X))W^O. \quad (2)$$

139 In addition, each Transformer block incorporates a residual connection (He et al., 2016), layer nor-  
 140 malization (Ba et al., 2016) and a position-wise feed-forward network (FFN).

141 **Parameter-Efficient Fine-Tuning.** We focus on the fine-tuning stage, assuming a Transformer-  
 142 based model pre-trained using any auxiliary task. For a downstream task, we are given labeled data  
 143  $\mathcal{X} = (X_1, \dots, X_n)$  and corresponding labels  $\mathcal{Y} = (y_1, \dots, y_n)$  to adapt the model. A simple ap-  
 144 proach is to train only a final classification or regression head while freezing the backbone, which is  
 145 efficient but often suboptimal. Full fine-tuning of both encoder and head improves performance but  
 146 requires substantial compute and memory. A recent alternative, Low-Rank Adaptation (LoRA) (Hu  
 147 et al., 2022), introduces low-rank trainable matrices  $A$  and  $B$  while keeping the original weight  
 148 matrix frozen. Specifically, for a dense layer weight  $W \in \mathbb{R}^{d \times k}$ , LoRA replaces it with:

$$149 \quad W' = W + \frac{\alpha}{r} BA,$$

150 where  $r$  is the rank,  $\alpha$  a scaling factor, and  $B \in \mathbb{R}^{d \times r}$ ,  $A \in \mathbb{R}^{r \times k}$  are learned during fine-tuning.

151 **Problem Setup.** Without loss of generality, we consider a 1-layer Transformer-based model (TBM)  
 152 where all activation functions are assumed to be 1-Lipschitz, which is the case for most commonly  
 153 used activations. The input space is  $\mathcal{X} \in [0, 1]^{n \times d}$ , representing normalized data such as images.  
 154 Fine-tuning is performed using an  $L$ -smooth loss function  $\mathcal{L}$ , optimized via gradient descent. Let  
 155  $W_*$  denote the local optimum to which the model converges. For a learning rate  $\eta \leq \frac{1}{L}$ , the update  
 156 rule for layer  $i$  at step  $t$  is:

$$157 \quad W_{t+1}^{(i)} = W_t^{(i)} - \eta \nabla \mathcal{L}(W_t^{(i)}).$$

162 While we focus on gradient descent for clarity, the theoretical insights extend to other optimizers  
 163 using similar analysis. Thus, our setup reflects a modeling choice rather than a limiting assumption.  
 164

## 165 4 ON THE ROBUSTNESS OF LORA 166

167 In this section, we aim to theoretically understand the connection between LoRA finetuning and the  
 168 resulting adversarial robustness, taking the head-only finetuning as a basis for comparison. We start  
 169 by formalizing the concept of expected adversarial robustness and, consequently, derive theoretical  
 170 insights for both the head-only finetuning and the LoRA counterpart, showcasing the difference in  
 171 terms of adversarial robustness.  
 172

### 173 4.1 ADVERSARIAL ROBUSTNESS 174

175 In this work, we focus on evasion attacks (Biggio et al., 2013; Pitropakis et al., 2019), which consist  
 176 of attacking the model at test or inference time. We consider that this setting is more adapted to real-  
 177 world scenarios, where in the majority of cases, the final user/attacker only has access to the model  
 178 at inference time. In this direction, let's consider a trained classifier  $f : \mathcal{X} \rightarrow \mathcal{Y}$  and let  $x \in \mathcal{X}$   
 179 be an input with its associated label vectors  $y \in \mathcal{Y}$ , such that  $f(x) = y$ . The goal of an attacker is  
 180 to craft a small additional perturbation to the input, such as to generate a point  $\tilde{x}$  whose prediction  
 181  $f(\tilde{x})$  is different from the original one. We note that the generated adversarial perturbation should be  
 182 similar to the original input, and therefore, we need to consider a similarity budget  $\epsilon$ , together with  
 183 the corresponding distance. For our current study, we consider the  $\ell_2$  distance and consequently  
 184 define our attack neighborhood of our input  $x$  with respect to an attack budget  $\epsilon$  as:  
 185

$$\mathcal{B}(x, \epsilon) = \{\tilde{x} \in \mathcal{X} : \|x - \tilde{x}\| \leq \epsilon\}$$

186 Given the previous neighborhood, the attacker aims to find within that neighborhood the points  
 187 that not only satisfy the adversarial aim of flipping the classification but also result in the worst  
 188 prediction. In this direction, given a finetuning strategy  $\zeta$  which is applied to our considered pre-  
 189 trained model  $f$ , the *adversarial risk* can be formulated as follows:  
 190

$$\mathcal{R}_\epsilon[f, \zeta] = \mathbb{E}_{x \in \mathcal{D}_{\mathcal{X}}} \left[ \sup_{\tilde{x} \in \mathcal{B}(x, \epsilon)} d_{\mathcal{Y}}(\zeta_f(\tilde{x}), \zeta_f(x)) \right]. \quad (3)$$

191 with  $d_{\mathcal{Y}}$  being any defined distances in the measurable output  $\mathcal{Y}$ . In the current work, and similar  
 192 to the input space, we consider  $\ell_2$ -norm as our distance metric for the output space. Note that there  
 193 exists an equivalence in terms of norm, and therefore, this latter choice can easily be extended to  
 194 other norms and doesn't limit our provided insights in any direction.  
 195

196 From an adversarial defense perspective, the objective is to ensure that the previously introduced  
 197 risk remains small, implying that it's harder to find a perturbation within the considered budget  $\epsilon$ ,  
 198 and consequently that the model predictions are stable within that neighborhood, reflecting the ad-  
 199 versarial robustness of the model. We can formalize this notion for a finetuning strategy as follows:  
 200

201 **Definition 1** (Adversarial Robustness). *The finetuning strategy  $\zeta$  is said to be  $(\epsilon, \gamma)$ -robust if its  
 202 adversarial risk with respect to the classifier  $f$  satisfies:  $\mathcal{R}_\epsilon[f, \zeta] \leq \gamma$ .*

203 We note that we approach the theoretical analysis from an upper-bound perspective (denoted as  $\gamma$ ),  
 204 since it is hard to compute the exact adversarial risk value. Obviously, the smaller the upper-bound,  
 205 the more robust the model is expected to be, and therefore by comparing the two quantities, we can  
 206 have an idea about the performance of the two considered finetuning strategies.  
 207

### 208 4.2 ON THE ROBUSTNESS OF LOW-RANK ADAPTATION 209

210 Building on the formal framework introduced previously, we now analyze the adversarial robust-  
 211 ness of Low-Rank Adaptation (LoRA) in comparison to the standard head-only finetuning strategy.  
 212 While LoRA is widely recognized for its effectiveness in improving downstream performance, often  
 213 measured by clean accuracy, this gain comes from its ability to modify a larger subset of the model's  
 214 parameters, including those within internal Transformer components. In contrast, head-only fine-  
 215 tuning restricts adaptation to the final classification layer, preserving the backbone of the pre-trained

model. This difference in parameter access raises, consequently, a natural question: Does the increased expressivity provided by LoRA come at the cost of adversarial robustness? Typically, while LoRA allows for better task-specific adaptation, it may also expose the model to increased vulnerability under test-time perturbations. To study this trade-off systematically, we adopt the notion of expected adversarial risk defined earlier, and derive upper bounds for both finetuning strategies under the same theoretical problem setup. Specifically, we consider a pre-trained Transformer-based model (TBM) denoted by  $f$ , and analyze its behavior under both head-only and LoRA-based finetuning. We consider  $f$  as a one-layer Transformer block with  $H$  dot-product self-attention heads, following the structure and the notations outlined in Section 3.

**Proposition 1.** *Let  $f: \mathcal{X} \rightarrow \mathcal{Y}$  be a pre-trained TBM-based model following the problem setup.*

*The head-only finetuning strategy  $\zeta_f^{\text{Head-only}}$  is  $(\epsilon, \gamma)$ -robust, with:  $\gamma_{\text{Head-Only}} = \left(\frac{d}{d-1}\right)^2 C_1 C_2 \epsilon$ ,*

$$C_1 = 1 + \|W_O\| \sqrt{H} \max_h \left[ \|W^{V,h}\| \left[ \frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1 \right] \right], \quad C_2 = (1 + \|W_{FFN}\|) \|W_{out}\|.$$

Proposition 1 provides a concrete expression for the adversarial risk bound under head-only finetuning, which depends explicitly on the norms of the model’s attention and feedforward weights. In the following, we extend this analysis to the case of LoRA-based finetuning, applying the same theoretical approach in order to establish a basis for comparison between the two strategies. Specifically, we consider that the LoRA is only applied to the query (Q) and value (V) projection matrices of the attention mechanism, as in the original proposed work.

**Theorem 1.** *Let  $f: \mathcal{X} \rightarrow \mathcal{Y}$  be the pre-trained TBM-based model following the considered problem setup. For the LoRA-based finetuning strategy  $\zeta_f^{\text{LoRA}}$ , where the LoRA is only applied to the main*

*Transformer part, is  $(\epsilon, \gamma)$ -robust, with:  $\gamma_{\text{LoRA}} = \left(\frac{d}{d-1}\right)^2 C'_1 C_2 \epsilon$ , where:*

$$C'_1 = 1 + \|W_O\| \sqrt{H} \max_h \left( \left[ \|W^{V,h}\| + \frac{\alpha}{r} \|A^{V,h}\| \|B^{V,h}\| \right] \left[ \frac{4}{\sqrt{d/H}} (\|W^{Q,h}\| + \frac{\alpha}{r} \|A^{Q,h}\| \|B^{Q,h}\|) \|W^{K,h}\| + 1 \right] \right).$$

The derived upper bounds in Proposition 1 and Theorem 1 provide a comparative theoretical framework for evaluating the adversarial robustness of head-only finetuning versus LoRA. While both bounds scale linearly with the perturbation radius  $\epsilon$  and share a similar structural dependence on network norms, the LoRA bound introduces additional terms involving the norms of the low-rank adaptation matrices  $A$  and  $B$ , scaled by the factor  $\alpha/r$ . These terms effectively inflate the expected adversarial risk of the model when subject to input perturbations, yielding a looser (i. e., higher) upper bound on the adversarial risk  $\mathcal{R}_\epsilon[f, \zeta]$  under LoRA.

This difference is intuitive and arises from the core design of LoRA, which introduces learnable low-rank updates into the internal weight matrices, specifically within the query, key, and value projections of the self-attention mechanism. By modifying these internal components, LoRA increases the space of adaptable parameters, enhancing task-specific expressivity and improving clean accuracy. However, this also creates additional pathways through which input perturbations can affect the output, making the model more vulnerable to adversarial attacks. In contrast, head-only finetuning restricts adaptation to the final classification layer, leaving the backbone unchanged and preserving the stability of the pre-trained representations. We consider that these results can also be categorized on the general robustness-performance trade-off, where, by aiming to have better performance, the model’s boundaries are adapted to the task, resulting in richer task-specific adaptation, but at the cost of amplifying the model’s response to small input perturbations. From a theoretical standpoint, this trade-off is captured directly by the looser robustness bound. Practically, it suggests that while LoRA may be preferable when downstream accuracy is the sole objective, it may lead to weaker performance in adversarial settings where robustness is critical. Specifically, for a final user, before choosing the right finetuning approach, an analysis of the objectives and the trade-off between clean and attacked accuracy should be done.

**Extension to Multi-Layer Transformers.** Although our theoretical analysis focuses on a single-layer Transformer-based model  $f$ , the results naturally extend to the multi-layer case. Specifically, a Transformer model with  $L$  layers, denoted as  $f^{(L)}$ , can be expressed as a composition of  $L$  single-layer functions:  $f^{(L)}(x) = f^{(L-1)} \circ f^{(L-2)} \circ \dots \circ f^{(1)}(x)$ . Under this formulation, and following

standard results from Lipschitz continuity, the overall adversarial risk bound  $\gamma$  for either finetuning strategy becomes a multiplicative composition of the bounds for each individual layer. That is, the robustness bound compounds across layers, maintaining the same structural form as in the single-layer case. We additionally note that the underlying assumptions of our problem setup (Section 3) still hold in the multi-layer setting. Since each layer operates on bounded activations (as we assume  $\mathcal{X} \subseteq [0, 1]^{n \times d}$ ), the composition of bounded functions preserves theoretical soundness. As a result, our robustness framework remains applicable in deeper architectures.

## 5 CONNECTING INITIALIZATION TO LORA’S ROBUSTNESS

In the previous section, we theoretically established a notable gap in adversarial robustness between head-only and LoRA-based finetuning strategies. Specifically, our analysis showed that LoRA exhibits a higher expected adversarial risk, suggesting reduced robustness to test-time perturbations. Motivated by this observation, we now turn to investigating whether this robustness gap can be influenced, or potentially mitigated and reduced, by specific choices in LoRA’s hyperparameters. In particular, we focus on how the *initialization of the low-rank matrices  $A$  and  $B$* , which define the LoRA updates, impacts adversarial robustness.

The effect of initialization (Hayou et al., 2024a) in LoRA has recently received increased attention. By design, the product of the two matrices is intended to start at zero to ensure that the training starts from the original weights of the pre-trained model. However, recent empirical studies (Hayou et al., 2024a) suggest that initializing  $A$  with random values and setting  $B$  to zero tends to yield better generalization and downstream performance than the reverse configuration. While these findings pertain to clean accuracy, their implications for adversarial robustness remain underexplored. In this section, we extend this line of inquiry by studying how the randomness in the initialization distribution of matrix  $A$ , which governs the initial adaptation direction, affects the final adversarial robustness of the finetuned model. Our goal is to understand whether certain initialization choices introduce more sensitivity to adversarial perturbations, and whether controlling the variance or structure of this randomness can lead to more robust LoRA configurations.

In this perspective, we consider the same setting as the one studied in the previous section, where  $f$  is a 1-Layer Transformer-based model and the aim is to link the initial weights with the resulting upper-bound on the expected adversarial robustness.

**Theorem 2.** *Let  $f: \mathcal{X} \rightarrow \mathcal{Y}$  be our pre-trained TBM-based model. Let’s consider the LoRA finetuning strategy, where all the low-rank matrices  $A$  in layer  $h$  are initialized as  $A_0^{Q,h}$  (for Query) and  $A_0^{V,h}$  (for Values), then the resulting  $C'_1$  constant in  $\gamma_{LoRA}$  ( Theorem 1) can be written as:*

$$C'_1 \leq K_1 (1 + \eta L)^t \max_h \|A_0^{(V,h)}\| + K_2 (1 + \eta L)^{2t} \max_h \|A_0^{(V,h)}\| \|A_0^{(Q,h)}\| + C,$$



with  $K_1, K_2$  and  $C$  being constants depending on the final weight norms (derived in Equation 12).

We observe that the upper bound derived in Theorem 2 directly links the norm of the chosen initialization matrix to the constant  $C'_1$ , which in turn influences  $\gamma_{LoRA}$  and thereby the model’s adversarial robustness. This result highlights that initialization, often treated as a secondary detail, plays a critical role in shaping LoRA’s robustness characteristics and should be carefully designed. Since the initialization also affects the model’s downstream performance, finding an appropriate trade-off between robustness and clean accuracy becomes crucial. In particular, the initialization of  $A$  should be designed to balance these objectives, enabling the construction of LoRA-based models that are both performant and robust. To better showcase the practical aspect of our theoretical result, we consider a practical application where we consider that the matrices are initialized from a Uniform distribution  $\mathcal{U}(-a, a)$ , where  $a$  is a parameter.

**Lemma 1.** *Consider LoRA matrices, with rank  $r$ , and for each head  $h = 1, \dots, H$  initialized with entries drawn i.i.d. from  $\mathcal{U}(-a, a)$  independently across heads and stacks. Then the expected value of the robustness constant  $C'_1$ , derived in Theorem 1, satisfies:*

$$\mathbb{E}[C'_1] = \mathcal{O}\left((1 + \eta L)^{2t} a^2 \left((\sqrt{r} + \sqrt{k}) + \sqrt{\log H}\right)^2\right).$$

The result of Lemma 1 establishes a direct relationship between the initialization parameter  $a$  of the Uniform distribution and the resulting upper bound on the expected adversarial robustness. Although



the analysis focuses on the Uniform case, a similar adaptation of Theorem 2, and similar reasoning can be extended to other initialization distributions.

## 6 EMPIRICAL EVALUATION

We empirically validate our theoretical insights using standard adversarial attacks across two widely used modalities: images and text. We start by outlining the experimental setup for both domains.

**Computer-Vision.** We have chosen to operate under two mainly widely used models, namely the Vision Transformer (ViT) (Dosovitskiy et al., 2021), which was the basis of our theoretical study, and the SwiN Transformer (Liu et al., 2021). For both models, we have considered the two mainly used adversarial attacks in this domain, which are the Fast Gradient Sign Method (FGSM) and the Proximal Gradient Descent (PGD), focusing on image classification using the CIFAR-10, CIFAR-100 (Krizhevsky et al., 2009), and ImageNet-100 (Russakovsky et al., 2015).

**Natural Language Processing (NLP).** We have chosen to operate through a number of different models, namely Bert-Base (Devlin et al., 2019), DistilBert-Base (Sanh et al., 2019), GPT2 (Radford et al., 2019), Gemma-2B (Team et al., 2024), llama3\_2\_1B (Dubey et al., 2024), Tiny-Llama (Zhang et al., 2024), and Mistral-7B (Jiang et al., 2023). For all models, we perform TextFooler (Jin et al., 2020) attack, and for some models, we also perform A2T (Yoo & Qi, 2021) attack. We focus the evaluation on the text classification task using IMdb (Maas et al., 2011), SST-2 (Socher et al., 2013), and Yelp Polarity (Zhang et al., 2015) datasets.

**Considered Metrics.** For both modalities, we report the clean/attacked accuracy and the success rate, which is the number of samples that were successfully attacked, meaning that the attack was successful in finding a perturbation within the budget that was able to flip the original classification.

We note that for all the models, the LoRA adaptation is applied to the whole self-attention component. The code to reproduce our results and experiments is provided in the Supplementary Materials, and additional details about the hyperparameters and the problem setup are provided in Appendix E.

378 6.1 EXPERIMENTAL RESULTS  
379

380 **Image-Based Evaluation.** Figure 2 (respectively Figure 8 in Appendix D.1) presents the aver-  
381 age success rate, and the corresponding standard deviations, of the FGSM attack for both head-  
382 only and LoRA finetuning strategies on a ViT (and SwiN, respectively), evaluated across multiple  
383 datasets and perturbation budgets. The empirical results align with our theoretical findings: across  
384 all datasets, LoRA consistently yields a higher attack success rate, indicating lower adversarial ro-  
385 bustness compared to head-only finetuning. Notably, the robustness gap between the two strategies  
386 can be substantial. For example, on ImageNet, the difference in attack success rate can reach up to  
387 20%, despite a clean accuracy gap of only around 1%. This performance contrast highlights that  
388 even small gains in clean performance under LoRA may come at a significant cost in adversarial  
389 settings. Similar insights are observed for the CIFAR dataset family, where by aiming for a small  
390 increase in clean robustness (around 3 – 4%), the resulting success rate can reach around 12 – 15%.

391  
392 Table 1: Average Clean Accuracy and Success Rate ( $\pm$  standard deviation) of the ViT and SwiN for  
393 both head-only and LoRA finetuning subject to FGSM and PGD attack on different datasets.

394 Model	395 Dataset	396 Strategy	397 Clean Accuracy $\uparrow$	398 Success Rate (FGSM) $\downarrow$	399 Success Rate (PGD) $\downarrow$
396 ViT	ImageNet	Head-Only	88.7 $\pm$ 0.2	28.9 $\pm$ 0.7	84.2 $\pm$ 0.4
		LoRA	88.9 $\pm$ 0.1	41.9 $\pm$ 0.8	96.4 $\pm$ 0.3
	CIFAR-10	Head-Only	97.4 $\pm$ 0.1	22.5 $\pm$ 0.3	88.4 $\pm$ 0.6
		LoRA	98.6 $\pm$ 0.1	27.1 $\pm$ 0.6	93.1 $\pm$ 0.3
	CIFAR-100	Head-Only	87.9 $\pm$ 0.9	42.9 $\pm$ 0.8	92.4 $\pm$ 0.7
		LoRA	90.8 $\pm$ 0.2	53.2 $\pm$ 0.4	96.2 $\pm$ 0.8
401 SwiN	ImageNet	Head-Only	89.8 $\pm$ 0.2	29.9 $\pm$ 0.8	90.2 $\pm$ 0.4
		LoRA	90.3 $\pm$ 0.1	36.6 $\pm$ 0.6	94.8 $\pm$ 0.7
	CIFAR-10	Head-Only	97.8 $\pm$ 0.1	26.1 $\pm$ 0.8	93.4 $\pm$ 0.2
		LoRA	98.5 $\pm$ 0.1	28.9 $\pm$ 0.7	95.1 $\pm$ 0.4
	CIFAR-100	Head-Only	87.6 $\pm$ 0.1	45.3 $\pm$ 0.6	94.2 $\pm$ 0.3
		LoRA	92.1 $\pm$ 0.3	50.2 $\pm$ 0.4	97.4 $\pm$ 0.2

406 **Text-Based Evaluation.** Figure 3 (and Figure 13 - Appendix D.3) presents the average success  
407 rates, along with standard deviations, for the TextFooler and A2T adversarial attacks applied to  
408 models fine-tuned using either head-only or LoRA strategies across varying perturbation budgets.  
409 The results in the NLP setting closely mirror the trends observed in computer vision, further rein-  
410 forcing the generality of our theoretical findings across modalities. In addition, Table 2 summarizes  
411 both clean accuracy and attack success rates under a fixed perturbation budget of 3 word substi-  
412 tutions. Across all evaluated models, LoRA fine-tuning generally shows lower robustness compared  
413 to head-only tuning. Additional experiments on other architectures and models are provided in  
414 Appendix D.3.

415  
416 Table 2: Average Clean Accuracy and Success Rate ( $\pm$  standard deviation) of BERT, DistilBERT  
417 and GPT-2 for head-only and LoRA finetuning subject to TextFooler and A2T on different datasets.

417 Model	418 Dataset	419 Strategy	420 Clean Accuracy $\uparrow$	421 Success Rate (TextFooler) $\downarrow$	422 Success Rate (A2T) $\downarrow$
420 BERT	IMDb	Head-Only	83.2 $\pm$ 0.9	11.3 $\pm$ 0.1	8.6 $\pm$ 0.6
		LoRA	90.5 $\pm$ 0.9	16.7 $\pm$ 1.0	14.9 $\pm$ 0.1
	SST-2	Head-Only	83.4 $\pm$ 0.5	54.0 $\pm$ 0.0	28.1 $\pm$ 0.7
		LoRA	92.1 $\pm$ 0.2	53.4 $\pm$ 1.6	22.9 $\pm$ 0.6
	Yelp Polarity	Head-Only	86.1 $\pm$ 1.5	14.7 $\pm$ 2.2	11.0 $\pm$ 1.1
		LoRA	92.6 $\pm$ 0.5	15.9 $\pm$ 0.8	11.5 $\pm$ 0.7
425 GPT-2	IMDb	Head-Only	85.7 $\pm$ 1.0	6.7 $\pm$ 0.8	11.0 $\pm$ 2.0
		LoRA	91.9 $\pm$ 0.9	7.9 $\pm$ 0.4	13.4 $\pm$ 2.3
	SST-2	Head-Only	82.1 $\pm$ 0.5	47.2 $\pm$ 0.9	30.1 $\pm$ 0.4
		LoRA	91.0 $\pm$ 1.3	54.6 $\pm$ 2.1	22.0 $\pm$ 0.8
	Yelp Polarity	Head-Only	85.2 $\pm$ 1.6	10.1 $\pm$ 0.8	12.3 $\pm$ 2.3
		LoRA	92.4 $\pm$ 0.5	7.9 $\pm$ 0.8	8.9 $\pm$ 1.4

429  
430 These findings underscore the practical significance of our theoretical analysis. While LoRA im-  
431 proves downstream performance in terms of clean accuracy, it also introduces increased vulnerabil-  
432 ity to adversarial perturbations. A key insight is that the gains in clean accuracy offered by LoRA

432 come at an adversarial cost, with performance degrading more severely under attack. This observation  
 433 highlights an important trade-off between clean accuracy and robustness, particularly in safety-  
 434 critical applications where reliability under distribution shift or adversarial threat is paramount. In  
 435 such contexts, clean accuracy alone is an insufficient metric and must be complemented by robust-  
 436 ness evaluations.

437

438 

## 6.2 EFFECT OF HYPER-PARAMETERS

439

440 

### 6.2.1 EFFECT OF INITIALIZATION

441

442 We further investigate the impact of initialization  
 443 strategies to demonstrate the practical relevance of  
 444 the theoretical insights from Section 5, particularly  
 445 Theorem 2. To this end, we evaluate several initial-  
 446 ization schemes for the LoRA matrix  $A$ . Specifi-  
 447 cally, we consider the default *Kaiming* initialization  
 448 used in the PEFT package, the well-known *Xavier*  
 449 initialization, and three additional classical distri-  
 450 butions: *Gaussian*, *Orthogonal*, and *Uniform*.

451

452 Figure 4 reports the average clean and attacked ac-  
 453 curacies across various adversarial budgets and ini-  
 454 tialization distributions. As anticipated, the choice  
 455 of initialization significantly influences the final ad-  
 456 versarial robustness. Although all distributions yield  
 457 similar clean accuracies (within a 2% range), the  
 458 attacked accuracies show a gap of up to 10% be-  
 459 tween the most and least robust initializations (Uni-  
 460 form versus Kaiming). These results indicate that se-  
 461 lecting an appropriate initialization can substantially  
 462 enhance robustness without sacrificing clean performance, thereby reducing the robustness gap be-  
 463 tween the LoRA and head-only finetuning strategies. Note that the additional results for the other  
 464 datasets are provided in Figure 11 (Appendix D.1). We further examine how the choice of initial  
 465 weight norm influences model vulnerability and adversarial robustness. Figure 5 reports results ob-  
 466 tained by varying the scaling factors of both Kaiming and Orthogonal initialization schemes. For  
 467 Kaiming initialization, note that the scaling factor is inversely related to the resulting weight norm  
 468 (via  $\sqrt{2}/(1 + a^2)$ ). Consistent with our theory, increasing the scaling parameter leads to larger  
 469 initial weight norms, which in turn produce higher  $\gamma$  values and reduced adversarial robustness.

470




Figure 4: Effect of the chosen initialization distribution on the resulting Attacked Accuracy of ImageNet.

471

472

473

474

475

476

477

478

479

480

481

482

483

484

485

486

487

488

489

490

491

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534

535

536

537

538

539

540

541

542

543

544

545

546

547

548

549

550

551

552

553

554

555

556

557

558

559

560

561

562

563

564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707

708

709

710

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810

811

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

1460

1461

1462

1463

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

1512

1513

1514

1515

1516

1517

1518

1519

1520

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

1534

1535

1536

1537

1538

1539

1540

1541

1542

1543

1544

1545

1546

1547

1548

1549

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

1565

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

1583

1584

1585

1586

1587

1588

1589

1590

1591

1592

1593

1594

1595

1596

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

1620

1621

1622

1623

1624

1625

1626

1627

1628

1629

1630

1631

1632

1633

1634

1635

1636

1637

1638

1639

1640

1641

1642

1643

Figure 6 presents the average attack success rates (with standard deviation) on CIFAR-10 and CIFAR-100 for different values of  $\alpha$ . Consistent with the theoretical insights that larger values of  $\alpha$  increase the upper bound, the empirical results confirm that increasing  $\alpha$  leads to higher attack success rates and thus reduced robustness. Interestingly, the widely adopted practice of setting  $\alpha = r$  appears suboptimal. Instead, using a smaller value such as  $\alpha = 1$  yields a reduction of approximately 10% in success rate, enhancing robustness and narrowing the gap with head-only fine-tuning.

### 6.3 GENERALIZING TO OTHER LORA ADAPTATIONS

Beyond the original LoRA formulation introduced by Hu et al. (2022), several extensions have been proposed to advance parameter-efficient fine-tuning. Our theoretical analysis focuses on this canonical variant to isolate how its parameterization influences adversarial robustness, but the resulting upper bound also suggests that certain adaptations may confer additional robustness benefits. For instance, DeLoRA Bini et al. (2025) and NB-LoRA

Wang et al. (2025) explicitly constrain the norm of the learned updates, which, when combined with our bound, implies a secondary robustness effect arising from tighter control of update magnitudes. Likewise, QLoRA Dettmers et al. (2023) employs 4-bit quantization to reduce memory usage while preserving downstream performance, and such quantization may incidentally lead to improved robustness. To evaluate these hypotheses, we empirically compare the adversarial robustness of standard LoRA with several of these adaptations. Finally, to illustrate how other PEFT strategies behave under adversarial perturbations, we also examine Prompt Tuning, in which a small set of learned continuous embeddings added to the input rather than modifying model weights, serves as the only trainable component. Figure 7 provides the empirical findings. As expected from our theoretical study, norm-bounded variants achieve higher robustness than standard LoRA, yet they still fall short of the robustness achieved by head-only fine-tuning. Additional results for NLP are provided in Figure 14.

## 7 CONCLUSION

In this work, we present the first theoretical analysis that explores the connection between LoRA as a fine-tuning strategy and the adversarial robustness of the resulting model. Our theoretical findings, supported by empirical results, indicate that the gains in clean accuracy achieved through LoRA come at the cost of increased vulnerability to adversarial attacks, particularly when compared to head-only fine-tuning. However, our analysis also highlights the important role of hyperparameters, specifically the scaling factor  $\alpha$  and the initialization scheme, in shaping this trade-off. We show that appropriate choices of these parameters can significantly reduce the robustness gap, yielding a more favorable balance between clean and attacked accuracy *without introducing additional constraints or computational overhead*, effectively offering a “free-lunch” improvement.

**Limitations.** While our work focuses on offering theoretical guidance for tuning LoRA’s hyperparameters, we believe it opens a new direction for designing LoRA variants that are not only effective on downstream tasks but also inherently more robust to adversarial perturbations.




Figure 6: Effect of the LoRA parameter  $\alpha$  on the resulting attack success rate.




Figure 7: Effect of the LoRA parameter  $\alpha$  on the resulting attack success rate.

540  
541  
**ETHICS STATEMENT**

542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
In this paper we study the adversarial robustness of LoRA-based models using only openly available datasets and pretrained models. Our work does not involve human subjects and therefore does not require IRB approval. All datasets used are publicly available and appropriately licensed. Although adversarial attacks are employed, they are standard, publicly available methods used solely to evaluate and improve model robustness. We believe that examining how these models respond to adversarial inputs is an important part of responsible AI research. By highlighting potential weaknesses, this line of work can help the community build systems that are more reliable, secure, and less vulnerable to misuse. In this context, to the best of our knowledge, this research does not raise ethical concerns related to discrimination, bias, privacy, or security. No conflicts of interest or legal compliance issues are associated with this work. We additionally note that LLMs were used only to assist with text refinement.

553  
554  
**REPRODUCIBILITY STATEMENT**  
555

556  
557  
558  
559  
560  
561  
562  
563  
We have made an effort to ensure that our results can be reproduced by others. All datasets and pretrained models we use are publicly available and are clearly referenced in the paper. The experimental setup, including how LoRA models are fine-tuned and how adversarial evaluations are carried out, is described in detail in the main text and the appendix (mainly Appendix E). The considered theoretical problem setup is clearly explained in Section 3 and all the theorem’s proofs and extended results are included in the appendix. Finally, to support independent verification, the code to reproduce our results is included in the Supplementary Materials and shall be made public upon publication.

564  
565  
**REFERENCES**  
566

567 Naveed Akhtar, Ajmal Mian, Navid Kardan, and Mubarak Shah. Advances in adversarial attacks  
568 and defenses in computer vision: A survey. *IEEE Access*, 9:155161–155196, 2021.

569 Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint*  
570 *arXiv:1607.06450*, 2016.

571 Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Gior-  
572 gio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Ma-  
573 chine learning and knowledge discovery in databases: European conference, ECML pKDD 2013,  
574 prague, czech Republic, September 23-27, 2013, proceedings, part III 13*, pp. 387–402. Springer,  
575 2013.

576 Massimo Bini, Leander Girrbach, and Zeynep Akata. Decoupling angles and strength in low-rank  
577 adaptation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL  
578 <https://openreview.net/forum?id=X1U74IwuxG>.

579 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for  
580 contrastive learning of visual representations. In *Proceedings of the 37th International Conference  
581 on Machine Learning*, ICML’20. JMLR.org, 2020.

582 Joana C Costa, Tiago Roxo, Hugo Proen  a, and Pedro RM In  cio. How deep learning sees the  
583 world: A survey on adversarial attacks & defenses. *IEEE Access*, 2024.

584 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning  
585 of quantized llms. *Advances in neural information processing systems*, 36:10088–10115, 2023.

586 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep  
587 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of  
588 the North American chapter of the association for computational linguistics: human language  
589 technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

594 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
 595 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit,  
 596 and Neil Houlsby. An image is worth 16x16 words: Transformers for image recogni-  
 597 tion at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.

598

599 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
 600 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.  
 601 *arXiv e-prints*, pp. arXiv–2407, 2024.

602

603 Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial  
 604 examples. In *International Conference on Learning Representations (ICLR)*, 2015.

605

606 Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. Parameter-efficient fine-tuning  
 607 for large models: A comprehensive survey. *Transactions on Machine Learning Research*, 2024.  
 608 ISSN 2835-8856. URL <https://openreview.net/forum?id=lIsCS8b6zj>.

609

610 Soufiane Hayou, Nikhil Ghosh, and Bin Yu. The impact of initialization on lora finetuning dynamics.  
 611 *Advances in Neural Information Processing Systems*, 37:117015–117040, 2024a.

612

613 Soufiane Hayou, Nikhil Ghosh, and Bin Yu. LoRA+: Efficient low rank adaptation of large  
 614 models. In *Forty-first International Conference on Machine Learning*, 2024b. URL <https://openreview.net/forum?id=NEv8YqBROO>.

615

616 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
 617 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
 618 770–778, 2016.

619

620 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 621 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

622

623 Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chap-  
 624 lot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier,  
 625 Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril,  
 626 Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL <https://arxiv.org/abs/2310.06825>.

627

628 Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. Is bert really robust? a strong baseline  
 629 for natural language attack on text classification and entailment. In *Proceedings of the AAAI  
 conference on artificial intelligence*, volume 34, pp. 8018–8025, 2020.

630

631 Damjan Kalajdzievski. A rank stabilization scaling factor for fine-tuning with lora. *arXiv preprint  
 arXiv:2312.03732*, 2023.

632

633 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint  
 arXiv:1412.6980*, 2014.

634

635 Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki M Asano. VeRA: Vector-based random matrix  
 636 adaptation. In *The Twelfth International Conference on Learning Representations*, 2024. URL  
 637 <https://openreview.net/forum?id=NjNfLdxr3A>.

638

639 Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In  
 640 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2661–  
 641 2671, 2019.

642

643 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.  
 644 2009.

645

646 Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo  
 647 Zhao. Loftq: LoRA-fine-tuning-aware quantization for large language models. In *The Twelfth  
 648 International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=LzPWWPAdy4>.

648 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-  
 649 Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. In *Forty-first*  
 650 *International Conference on Machine Learning*, 2024.

651 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.  
 652 Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the*  
 653 *IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.

654 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.

655 Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts.  
 656 Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the*  
 657 *association for computational linguistics: Human language technologies*, pp. 142–150, 2011.

658 Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu.  
 659 Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*,  
 660 2017.

661 Nikolaos Pitropakis, Emmanouil Panaousis, Thanassis Giannetsos, Eleftherios Anastasiadis, and  
 662 George Loukas. A taxonomy and survey of attacks against machine learning. *Computer Science*  
 663 *Review*, 34:100199, 2019.

664 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language  
 665 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

666 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng  
 667 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual  
 668 recognition challenge. *International journal of computer vision*, 115:211–252, 2015.

669 Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of  
 670 bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.

671 Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng,  
 672 and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment  
 673 treebank. In *Proceedings of the 2013 conference on empirical methods in natural language pro-*  
 674 *cessing*, pp. 1631–1642, 2013.

675 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhu-  
 676 patiraju, Léonard Hussonot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma  
 677 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.

678 Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks  
 679 to adversarial example defenses. *Advances in neural information processing systems*, 33:1633–  
 680 1645, 2020.

681 Bohdan Turbal, Anastasiia Mazur, Jiaxu Zhao, and Mykola Pechenizkiy. On adversarial robust-  
 682 ness of language models in transfer learning. In *Workshop on Socially Responsible Language*  
 683 *Modelling Research*, 2024. URL <https://openreview.net/forum?id=DUMVB9a9sX>.

684 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
 685 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*  
 686 *tion processing systems*, 30, 2017.

687 Ruigang Wang, Krishnamurthy Dvijotham, and Ian R Manchester. Norm-bounded low-rank adap-  
 688 *ation*. *arXiv preprint arXiv:2501.19050*, 2025.

689 Hao Wu, Xiangfeng Luo, Jianqi Gao, and Dian Huang. Improving text processing via adversarial  
 690 low-rank adaptation. *Machine Learning*, 114(9):196, 2025.

691 Xilie Xu, Jingfeng Zhang, and Mohan Kankanhalli. Autolora: An automated robust fine-tuning  
 692 framework. In *The Twelfth International Conference on Learning Representations*, 2024. URL  
 693 <https://openreview.net/forum?id=09xFexjhqE>.

702 Xu Yang, Chen Liu, and Ying Wei. Mixture of adversarial loras: Boosting robust generalization in  
703 meta-tuning. *Advances in Neural Information Processing Systems*, 37:115176–115207, 2024a.  
704

705 Zeyu Yang, Zhao Meng, Xiaochen Zheng, and Roger Wattendofer. Assessing adversarial robustness  
706 of large language models: An empirical study. *arXiv preprint arXiv:2405.02764*, 2024b.

707 Jin Yong Yoo and Yanjun Qi. Towards improving adversarial training of NLP models. In Marie-  
708 Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Findings of the*  
709 *Association for Computational Linguistics: EMNLP 2021*, pp. 945–956, Punta Cana, Dominican  
710 Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.  
711 *findings-emnlp.81*. URL [https://aclanthology.org/2021.findings-emnlp.](https://aclanthology.org/2021.findings-emnlp.81/)  
712 81/.

713 Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. Tinyllama: An open-source small  
714 language model. *arXiv preprint arXiv:2401.02385*, 2024.

715

716 Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text clas-  
717 sification. *Advances in neural information processing systems*, 28, 2015.

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756  
757  
758  
759  
760  
761

## Supplementary Material

### A PROOF OF PROPOSITION 1

**Proposition.** Let  $f: \mathcal{X} \rightarrow \mathcal{Y}$  be the pre-trained TBM-based model following the considered problem setup. The head-only finetuning strategy  $\zeta_f^{\text{Head-only}}$  is  $(\epsilon, \gamma_{\text{Head-Only}})$ -robust, with:

$$\gamma_{\text{Head-Only}} = \left( \frac{d}{d-1} \right)^2 C_1 C_2 \epsilon,$$

with:  $C_1 = (1 + \|W_O\| \sqrt{H} \max_h [\|W^{V,h}\| [\frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1]]),$

$$C_2 = (1 + \|W_{FFN}\|) \|W_{out}\|,$$

and  $W$  being the different weights of the models (as explained in Section 3).

*Proof.* Let's consider our input  $X \in \mathcal{X}$  composed of  $n$  tokens  $x_i \in \mathbb{R}^d$ . We consider that our model  $f$  is built using the dot-product self-attention as referred to in Equation 1 and reformulated as:

$$\begin{aligned} \text{AH}(x) &= \text{Softmax}\left(\frac{(XW^Q)(XW^K)^T}{\sqrt{\frac{D}{H}}}\right)XW^V \\ &= PXW^V = h(X)W^V, \end{aligned}$$

where  $W^Q, W^K, W^V$  are learnable weights of the model. Let's consider the function  $h(X)$ , we can write:

$$f(X) = PX = \text{Softmax}(XA^T X^T)X$$

$$f(X) = PX = \text{Softmax}(XA^T X^T)X = \begin{bmatrix} h_1(X)^\top \\ \vdots \\ h_n(X)^\top \end{bmatrix} \in \mathbb{R}^{n \times d}, \quad \text{with:}$$

$$A = \frac{W^K W^{Q^\top}}{\sqrt{d/H}} \in \mathbb{R}^{d \times d} \quad \text{and} \quad h_i(X) = \sum_{j=1}^n P_{ij} x_j \quad \text{with} \quad P_i^\top = \text{Softmax}(XAx_i).$$

By analyzing the partial derivatives, we can directly write the following regarding the Jacobian matrix of  $h$ :

$$J_{ij} = X^\top P^{(i)} E_{ji} X A^\top + \delta_{ij} (X^\top P^{(i)} X A) + P_{ij} I_d,$$

with:

- $P^{(i)} = \text{diag}(P_{i:}) - P_{i:}^\top P_{i:}$ , [Softmax derivate]
- $E_{ji}$  is the  $(n \times n)$  matrix with a single 1 in position  $(j, i)$ .

Based on this, two elements arises:

$$\text{If } i \neq j, \quad J_{ij} = X^\top P^{(i)} E_{ji} X A^\top + P_{ij} I, \quad (4)$$

$$\text{If } i = j, \quad J_{ii} = X^\top P^{(i)} E_{ii} X A^\top + X^\top P^{(i)} X A + P_{ii} I. \quad (5)$$

810 We recall that the input images are considered to be normalized, and therefore we can write:  
 811  
 812

$$\|X\| \leq 1$$

813 Additionally, since  $P_{i:}$  is the output of the softmax, then can be considered a probability distribution.  
 814 Therefore,  $\sigma_{\max}(\text{diag}(p)) \leq 1$  and  $pp^T$  has rank 1:  
 815

$$\|P^{(i)}\| = \|\text{diag}(P_{i:}) - P_{i:}^T P_{i:}\| \leq 2$$

816 **Case 1.** We start by considering the first case  $i \neq j$ , in which we have:  
 817

$$J_{ij} = X^T P^{(i)} E_{ji} X A^T + P_{ij} I.$$

818 Consequently we have the following:  
 819

$$\begin{aligned} \|J_{ij}\| &\leq \|X^T P^{(i)} E_{ji} X A^T\| + \|P_{ij} I\| \\ &\leq 2 \|A\| + 1 \\ &\leq \|A\| + 1 \end{aligned}$$

820 **Case 2.** For the second case  $i = j$ , we have the following:  
 821

$$J_{ii} = X^T P^{(i)} E_{ii} X A^T + X^T P^{(i)} X A + P_{ii} I.$$

822 We apply the same analogy as the previous case:  
 823

$$\begin{aligned} \|J_{ii}\| &\leq \|X^T P^{(i)} E_{ii} X A^T\| + \|X^T P^{(i)} X A\| + \|P_{ii} I\| \\ &\leq 2 \|A\| + 2 \|A\| + 1 \\ &\leq 4 \|A\| + 1 \end{aligned}$$

824 So overall, we have the following:  
 825

$$\|J_{ij}\|_{op} \leq \begin{cases} 2 \|A\| + 1, & \text{if } i \neq j, \\ 4 \|A\| + 1, & \text{if } i = j. \end{cases}$$

826 So with our theoretical assumptions, the Jacobian is bounded and we have:  $\mathcal{L}_h \leq 4 \|A\| + 1$ .  
 827

828 Specifically, for an attention head  $h$ , we have the following computation taking into account the  
 829 different learnable weights:  
 830

$$\mathcal{L}_{head} \leq \|W^{V,h}\| \left[ \frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1 \right]$$

831 Since  $f$  is represented by  $H$  separate attention head, then their concatenated output as explained in  
 832 Equation 2 is subject to the following:  
 833

$$\begin{aligned} \mathcal{L}_{MH} &\leq \|W_O\| \sqrt{H} \max_h [\mathcal{L}_{head}] \\ &\leq \|W_O\| \sqrt{H} \max_h [\|W^{V,h}\| \left[ \frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1 \right]] \end{aligned}$$

834 Finally, we need to take into account the application of the FFN and LN (with its parameters  $\gamma = 1$   
 835 and  $\beta = 1$ ). In addition, giving that we consider the head-only fine-tuning strategy, we consider that  
 836 a final linear layer  $W_{out}$  is trained. Since ReLU is 1-Lipschitz, we have the following result:  
 837

$$\begin{aligned} \mathcal{L}_f &\leq L_{LN}^2 (1 + \mathcal{L}_{MH}) (1 + L_{FFN}) \\ &\leq \left( \frac{d}{d-1} \right)^2 (1 + \mathcal{L}_{MH}) (1 + \|W_{FFN}\|) \\ &\leq \left( \frac{d}{d-1} \right)^2 (1 + \|W_O\| \sqrt{H} \max_h [\|W^{V,h}\| \left[ \frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1 \right]]) (1 + \|W_{FFN}\|) \\ &\leq \left( \frac{d}{d-1} \right)^2 A_1 A_2, \end{aligned}$$

864

865

$$866 \quad \text{with } A_1 = \left(1 + \|W_O\| \sqrt{H} \max_h \left[ \|W^{V,h}\| \left[ \frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1 \right] \right] \right)$$

867

$$868 \quad A_2 = \left(1 + \|W_{FFN}\|\right) \|W_{\text{out}}\|$$

869

870 Let's now consider a perturbed input  $\tilde{x} \in \mathcal{B}(x, \epsilon)$  as defined in Section 4.1. The previous upper-  
871 bound applies to any given point within that budget, and therefore we have:

872

873

$$874 \quad \sup_{\tilde{x} \in \mathcal{B}(x, \epsilon)} d_{\mathcal{Y}}(\zeta_f(\tilde{x}), \zeta_f(x)) \leq \mathcal{L}_f \epsilon$$

875

876

877 We can therefore conclude that, in respect of Definition 1, the head-only finetuning strategy is is  
878  $(\epsilon, \gamma_{\text{Head-Only}})$ -robust, with:

879

880

$$881 \quad \gamma_{\text{Head-Only}} = \left( \frac{d}{d-1} \right)^2 C_1 C_2 \epsilon,$$

882

883

$$884 \quad \text{with: } C_1 = \left(1 + \|W_O\| \sqrt{H} \max_h \left[ \|W^{V,h}\| \left[ \frac{4}{\sqrt{d/H}} \|W^{Q,h}\| \|W^{K,h}\| + 1 \right] \right] \right),$$

885

$$886 \quad C_2 = \left(1 + \|W_{FFN}\|\right) \|W_{\text{out}}\|,$$

887

888

## 889 B PROOF OF THEOREM 1

890

891

892 **Theorem.** Let  $f: \mathcal{X} \rightarrow \mathcal{Y}$  be the pre-trained TBM-based model following the considered problem  
893 setup. For the LoRA-based finetuning strategy  $\zeta_f^{\text{LoRA}}$ , where the LoRA is only applied to the main  
894 transformer part, is  $(\epsilon, \gamma_{\text{LoRA}})$ -robust, with:

895

896

$$897 \quad \gamma_{\text{LoRA}} = \left( \frac{d}{d-1} \right)^2 C'_1 C_2 \epsilon,$$

898

899

$$900 \quad \text{with: } C'_1 = 1 + \|W_O\| \sqrt{H} \max_h \left[ \|W^{V,h}\| + \frac{\alpha}{r} \|A^{V,h}\| \|B^{V,h}\| \right] \\ 901 \quad \left[ \frac{4}{\sqrt{d/H}} \left[ \|W^{Q,h}\| + \frac{\alpha}{r} \|A^{Q,h}\| \|B^{Q,h}\| \right] \|W^{K,h}\| + 1 \right]$$

902

903

904 *Proof.* In this part, we consider the LoRA-based finetuning. In this perspective, we follow the same  
905 analogy as the previous proof. Specifically, Let  $X \in \mathcal{X}$  be our input composed of  $n$  tokens  $x_i \in \mathbb{R}^d$ .  
906 We consider the same model  $f$  which is built using the dot-product self-attention as referred to in  
907 Equation 1 and reformulated as:

908

909

910

911

$$912 \quad \text{AH}(x) = \text{Softmax} \left( \frac{(XW^Q)(XW^K)^T}{\sqrt{\frac{D}{H}}} \right) XW^V \\ 913 \quad = PXW^V = h(X)W^V,$$

914

915

916

917 We recall that in the case of LoRA, two additional matrices  $A$  and  $B$  are learnable during the fine-  
918 tuning. Specifically, given a weight matrix  $W \in \mathbb{R}^{d \times k}$  in a model, it is substituted by the following:

919

920

921

922

923

924  $W' = W + \frac{\alpha}{r} BA,$   
925 where  $r$  is the rank of the low-rank adaptation,  $\alpha$  is the scaling factor, and  $B \in \mathbb{R}^{d \times r}$  and  $A \in \mathbb{R}^{r \times k}$   
926 are learnable weight matrices learned during the finetuning process.

918 From the previous section, we have the following Lipschitz bound for the head-only finetuning  
 919 strategy:  
 920

$$\begin{aligned}
 922 \quad \mathcal{L}'_f &\leq L_{LN}^2(1 + \mathcal{L}_{MH})(1 + L_{FFN}) \\
 923 \quad &\leq \left(\frac{d}{d-1}\right)^2(1 + \mathcal{L}_{MH})(1 + \|W_{FFN}\|) \\
 924 \quad &\leq \left(\frac{d}{d-1}\right)^2(1 + \|W_O\|\sqrt{H} \max_h [\|W'^{V,h}\| \left[\frac{4}{\sqrt{d/H}} \|W'^{Q,h}\| \|W'^{K,h}\| + 1\right]])(1 + \|W_{FFN}\|)
 \end{aligned}$$

928 We consider that the LoRA finetuning is only applied to the main core of Transformer, specifically  
 929 to the query  $Q$  and value  $V$  matrices of the attention mechanism as in the original work. We can  
 930 therefore continue the previous computation by including the corresponding values:  
 931

$$\begin{aligned}
 933 \quad \mathcal{L}'_f &\leq \left(\frac{d}{d-1}\right)^2(1 + \|W_O\|\sqrt{H} \max_h [\|W'^{V,h}\| \left[\frac{4}{\sqrt{d/H}} \|W'^{Q,h}\| \|W'^{K,h}\| + 1\right]])(1 + \|W_{FFN}\|) \\
 934 \quad &\leq \left(\frac{d}{d-1}\right)^2(1 + \|W_O\|\sqrt{H} \max_h [\left[\|W^{V,h}\| + \frac{\alpha}{r} \|A^{V,h}\| \|B^{V,h}\|\right] \\
 935 \quad &\quad \left[\frac{4}{\sqrt{d/H}} \left[\|W^{Q,h}\| + \frac{\alpha}{r} \|A^{Q,h}\| \|B^{Q,h}\|\right] \|W^{K,h}\| + 1\right]])(1 + \|W_{FFN}\|) \\
 936 \quad &\leq \left(\frac{d}{d-1}\right)^2 A'_1 A_2,
 \end{aligned}$$

942 Similar to the previous proof, let's consider a perturbed input  $\tilde{x} \in \mathcal{B}(x, \epsilon)$  as defined in Section 4.1.  
 943 The previous upper-bound applies to any given point within that budget, and therefore we have:  
 944

$$\sup_{\tilde{x} \in \mathcal{B}(x, \epsilon)} d_{\mathcal{Y}}(\zeta_f(\tilde{x}), \zeta_f(x)) \leq \mathcal{L}'_f \epsilon$$

945 We can therefore conclude that, in respect of Definition 1, the LoRA finetuning strategy is  $(\epsilon, \gamma_{\text{LoRA}})$ -  
 946 *robust*, with:  
 947

$$\begin{aligned}
 948 \quad \gamma_{\text{LoRA}} &= \left(\frac{d}{d-1}\right)^2 C'_1 C_2 \epsilon, \\
 949 \quad \text{with: } C'_1 &= 1 + \|W_O\|\sqrt{H} \max_h \left[\|W^{V,h}\| + \frac{\alpha}{r} \|A^{V,h}\| \|B^{V,h}\|\right] \\
 950 \quad &\quad \left[\frac{4}{\sqrt{d/H}} \left[\|W^{Q,h}\| + \frac{\alpha}{r} \|A^{Q,h}\| \|B^{Q,h}\|\right] \|W^{K,h}\| + 1\right]
 \end{aligned}$$

□

## C PROOF OF THEOREM 2

955 **Theorem.** Let  $f: \mathcal{X} \rightarrow \mathcal{Y}$  be our pre-trained TBM-based model. Let's consider the LoRA finetuning  
 956 strategy, where all the low-rank matrices  $A$  in layer  $h$  are initialized as  $A_0^{Q,h}$  (for Query) and  $A_0^{V,h}$   
 957 (for Values), then the resulting  $C'_1$  constant in  $\gamma_{\text{LoRA}}$  (Theorem 1) can be written as:  
 958

$$C'_1 \leq K_1 (1 + \eta L)^t \max_h \|A_0^{(V,h)}\| + K_2 (1 + \eta L)^{2t} \max_h \|A_0^{(V,h)}\| \|A_0^{(Q,h)}\| + C,$$

959 with  $K_1, K_2$  and  $C$  being the constants depending on the final weight norms (derived in Equation  
 960 12).  
 961

972 *Proof.* Let's now consider the effect of Initialization distribution on the final adversarial robustness  
 973 of the LoRA finetuning. Specifically, we consider the same settings as in prior work, where the  $B$   
 974 matrix is set to 0 and only the  $A$  matrix is initialized.

975 The gradient descent update at finetuning epoch  $t$  for our matrix  $A$  (at any layer) is written as:  
 976

$$977 \quad A_{t+1}^{(Q,h)} = A_t^{(Q,h)} - \eta \nabla \mathcal{L}(A_t^{(Q,h)}).$$

979 As specified in our problem setup in Section 3, we consider that our loss function  $\mathcal{L}$  to be  $L$ -smooth,  
 980 we can hence write the following result:  
 981

$$982 \quad \|\nabla \mathcal{L}(A_t^{(Q,h)})\| \leq L \|A_t^{(Q,h)} - A_*^{(Q,h)}\|.$$

984 Consequently, after  $t$  training epochs, we can write:  
 985

$$\begin{aligned} 986 \quad \|A_t^{(Q,h)}\| &= \|A_{t-1}^{(Q,h)} - \eta \nabla \mathcal{L}(A_{t-1}^{(Q,h)})\| \\ 987 &\leq \|A_{t-1}^{(Q,h)}\| + \eta L \|A_{t-1}^{(Q,h)} - A_*^{(Q,h)}\| \\ 988 &\leq (1 + \eta L) \|A_{t-1}^{(Q,h)}\| + \eta L \|A_*^{(Q,h)}\|. \end{aligned}$$

991 In addition, we suppose that the considered learning rate is chosen as  $\eta \leq \frac{1}{L}$ . Consequently, we can  
 992 write based on the previous formulation and by using recursion:  
 993

$$994 \quad \|A_t^{(Q,h)}\| \leq (1 + \eta L)^t \|A_0^{(Q,h)}\| + \sum_{h=0}^t 2^h \|A_*^{(Q,h)}\| \quad (6)$$

$$997 \quad \leq (1 + \eta L)^t \|A_0^{(Q,h)}\| + 2^{t+1} \|A_*^{(Q,h)}\| \quad (7)$$

$$999 \quad \leq (1 + \eta L)^t \|A_0^{(Q,h)}\| + 2^{t+1} \|A^{(Q,h)}\| \quad (8)$$

1001 *Remark.* We denote  $A_*$  (which are the converged final weights) as  $A$  directly to be inline with the  
 1002 previous theorems and results.  
 1003

1004 We additionally note that a similar analogy applies to the matrix  $B$ . Since we consider that specific  
 1005 matrix to be initialized to zero, the resulting terms of the initialization is therefore set to zero and  
 1006 only the final weight norm is seen in the upper-bound.

1007 Consequently, and from the previous part, we had that the LoRA finetuning is  $(\epsilon, \gamma_{\text{LoRA}})$ -robust,  
 1008 with:

$$1009 \quad \gamma_{\text{LoRA}} = \left( \frac{d}{d-1} \right)^2 C'_1 C_2 \epsilon,$$

$$\begin{aligned} 1013 \quad \text{with: } C'_1 &= 1 + \|W_O\| \sqrt{H} \max_h \left[ \|W^{V,h}\| + \frac{\alpha}{r} \|A^{V,h}\| \|B^{V,h}\| \right] \\ 1014 &\quad \left[ \frac{4}{\sqrt{d/H}} \left[ \|W^{Q,h}\| + \frac{\alpha}{r} \|A^{Q,h}\| \|B^{Q,h}\| \right] \|W^{K,h}\| + 1 \right] \end{aligned}$$

1018 Using the result from Equation 8, we can connect the previous resulting bound to the initialization,  
 1019 resulting in the following:  
 1020

$$\begin{aligned} 1022 \quad \text{with: } C'_1 &= 1 + \|W_O\| \sqrt{H} \max_h \left[ \|W^{V,h}\| + \frac{\alpha}{r} \|B^{V,h}\| \left[ (1 + \eta L)^t \|A_0^{(V,h)}\| + 2^{t+1} \|A^{(V,h)}\| \right] \right] \\ 1023 &\quad \left[ \frac{4}{\sqrt{d/H}} \left[ \|W^{Q,h}\| + \frac{\alpha}{r} \|B^{Q,h}\| \left[ (1 + \eta L)^t \|A_0^{(Q,h)}\| + 2^{t+1} \|A^{(Q,h)}\| \right] \right] \|W^{K,h}\| + 1 \right] \end{aligned}$$

1026  
1027

From the previous result, we can separate the two main terms, as follows:

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

$$C'_1 = 1 + \|W_O\| \sqrt{H} \max_h \left\{ \underbrace{\left[ \|W^{V,h}\| + \frac{\alpha}{r} \|B^{V,h}\| ((1 + \eta L)^t \|A_0^{(V,h)}\| + 2^{t+1} \|A^{(V,h)}\|) \right]}_{=a_h} \right\} \quad (9)$$

1032

1033

1034

1035

1036

1037

The previous two elements can be written as:

1038

1039

1040

1041

1042

1043

1044

$$\begin{aligned} a_h &\leq C_h^{(V)} + \frac{\alpha}{r} \|B^{V,h}\| (1 + \eta L)^t \|A_0^{(V,h)}\| \\ b_h &\leq C_h^{(QK)} + \frac{4}{\sqrt{d/H}} \frac{\alpha}{r} \|B^{Q,h}\| (1 + \eta L)^t \|A_0^{(Q,h)}\| \|W^{K,h}\| \end{aligned}$$

1044

Note that since everything is positive (given that the norms by definition are positive), then we have:

1045

1046

1047

1048

1049

1050

Consequently, we can write:

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

Let's define the following:  $A_0^Q = \max_h \|A_0^{(Q,h)}\|$ , and  $A_0^V = \max_h \|A_0^{(V,h)}\|$ .

Then we can finally write:

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

**Proof of Lemma 1:**

**Lemma.** Consider LoRA matrices for each head  $h = 1, \dots, H$  initialized with entries drawn i.i.d. from  $\mathcal{U}(-a, a)$  independently across heads and stacks. Then the expected value of the robustness constant  $C'_1$  satisfies

$$\mathbb{E}[C'_1] = \mathcal{O}\left((1 + \eta L)^{2t} a^2 \left((\sqrt{r} + \sqrt{k}) + \sqrt{\log H}\right)^2\right).$$

1080 *Proof.* From the Theorem2, for sufficiently large  $t$  we have  
 1081

$$1082 C'_1 = \mathcal{O}\left((1 + \eta L)^{2t} \max_h \|A_0^{(V,h)}\| \max_h \|A_0^{(Q,h)}\|\right). \quad (14)$$

1084 For a random matrix  $A \in \mathbb{R}^{r \times k}$  with i.i.d. entries from  $\mathcal{U}(-a, a)$ , matrix concentration bounds yield  
 1085

$$1086 \mathbb{E}[\|A\|] \leq C_1 a(\sqrt{r} + \sqrt{k}), \quad (15)$$

1088 for a universal constant  $C_1 > 0$ . Taking the maximum over  $H$  independent heads gives  
 1089

$$1090 \mathbb{E}\left[\max_{h=1, \dots, H} \|A^{(h)}\|\right] \leq C_2 a((\sqrt{r} + \sqrt{k}) + \sqrt{\log H}). \quad (16)$$

1092 By independence of Q and V stacks,  
 1093


$$1094 \mathbb{E}\left[\max_h \|A_0^{(V,h)}\| \max_h \|A_0^{(Q,h)}\|\right] = \mathbb{E}\left[\max_h \|A_0^{(V,h)}\|\right] \cdot \mathbb{E}\left[\max_h \|A_0^{(Q,h)}\|\right], \quad (17)$$

1096 which yields the stated bound after plugging in Equation 14, multiplication and absorbing constants.  $\square$   
 1097

## 1099 D ADDITIONAL RESULTS

### 1100 D.1 ADDITIONAL RESULTS CV

1103 **Results SwiN-based Transformer.** While in Section 6, Figure 2 provides the results of the success  
 1104 rate of the ViT, we additionally provide the results on the SwiN, which show similar tendencies and  
 1105 insights on the link between loRA and adversarial robustness compared to head-only finetuning.  
 1106



1119 Figure 8: Success Rate of the FGSM Attack when applied to a SwiN-based Model and different  
 1120 datasets and attack budgets.  
 1121

1122 **Clean/Attacked Accuracy.** Additionally, as explained in Section 6, we study the adversarial robustness  
 1123 from an attack success rate perspective as a representative metric, but we additionally report  
 1124 the attacked accuracy, which also has the same insights. Note that the success rate is simply a rep-  
 1125 resentative function of the attack accuracy (without taking into account the erroneous predictions  
 1126 of the models). In this perspective, Figure 9 provides the results of the average clean and attacked  
 1127 accuracy of a ViT when subject to the FGSM attack for different attack budgets, while Figure 10  
 1128 provides the results for the SwiN-based transformer.  
 1129

1130 **Other Results on Initialization.** In line with what was presented in Section 6.2, we extend the  
 1131 study to take into account the other considered datasets. In this perspective, Figure 11 provides the  
 1132 corresponding results on the ImageNet, the CIFAR-10 and CIFAR-100.  
 1133

### D.2 ON THE EFFECT OF RANK




Figure 9: Clean/Attacked Accuracy of a ViT when subject to the FGSM attack for different budgets and datasets.




Figure 10: Clean/Attacked Accuracy of a SwiN-based model when subject to the FGSM attack for different budgets and datasets.

In addition to studying the effect of the  $\alpha$  value in the empirical evaluation in Section 6 (Figure 6), we evaluate the effect of the chosen rank value  $r$  which is also present in the computed upper-bound. Figure 12 shows the resulting study, where we see that increasing the value of the rank  $r$  results in decreasing the attack success rate, in which is in accordance with our theoretical results, where  $\gamma$  is inversely proportional to the chosen value  $r$ .

### D.3 ADDITIONAL RESULTS NLP

**Results for A2T Attack.** While in Section 6, Figure 3 presents the results of TextFooler attack on head-only and LoRA finetuning, Figure 13 presents the attack results from A2T attack. Overall, we have similar overall observations for the adversarial robustness comparison between head-only and LoRA finetuning.

**Adversarial Robustness Comparison of Norm-Bounded LoRA.** While our theory show that LoRA is less robust than Head-Only finetuning, recent research work has also demonstrate that fine-tuning under bounded parameter norms improves the adversarial robustness of LoRA. Specifically, DeLoRA (Bini et al., 2025) normalizes and scales low-rank updates to decouple the direction of weight changes from their magnitude. More recently, NB-LoRA (Wang et al., 2025) reparameterizes low-rank adaptation matrices so that its singular values are explicitly bounded, yielding improved training stability and hyper-parameter robustness compared to LoRA. Figure 14 shows the empirical comparison between LoRA, DeLoRA, NB-LoRA and Head-Only fine-tuning. In




Figure 12: Effect of the rank value parameter  $r$  of LoRA on the resulting attack success rate.




Figure 11: Clean/Attacked Accuracy of a ViT when subject to FGSM under different initialization distribution.




Figure 13: Attack Success Rate of the A2T Attack on BERT, DistilBERT, and GPT-2 when applied to IMDb dataset with different attack budget (number of words changed).

general, we find that both DeLoRA and NB-LoRA improves the adversarial robustness of the model as compared to LoRA.

**Clean/Attacked Accuracy.** In addition to attack success rate, Table 3 and 4 and provides detailed results of clean accuracy and adversarial accuracy of TextFooler and A2T attack.

Table 3: Original accuracy (Orig Acc), attack success rate (ASR, %) and adversarial accuracy (Adv Acc, %) across datasets with TextFooler attack. Attack budget: max number of words changed is 3.

	Dataset Metric	IMDb		SST-2		Yelp Polarity			
		Orig Acc	ASR	Adv Acc	Orig Acc	ASR	Adv Acc	Orig Acc	ASR
<b>BERT</b>	Head-Only	83.20 $\pm$ 0.87	11.27 $\pm$ 0.12	71.93 $\pm$ 0.83	83.40 $\pm$ 0.53	54.00 $\pm$ 0.00	29.40 $\pm$ 0.53	86.13 $\pm$ 1.53	14.73 $\pm$ 2.16
	LoRA	90.53 $\pm$ 0.92	16.67 $\pm$ 1.01	73.87 $\pm$ 1.89	92.13 $\pm$ 0.23	53.40 $\pm$ 1.64	38.73 $\pm$ 1.86	92.60 $\pm$ 0.53	15.93 $\pm$ 0.81
<b>DistilBERT</b>	Head-Only	83.60 $\pm$ 0.40	11.40 $\pm$ 0.35	72.20 $\pm$ 0.72	82.13 $\pm$ 0.31	52.00 $\pm$ 0.72	30.13 $\pm$ 0.50	86.07 $\pm$ 1.86	14.60 $\pm$ 1.51
	LoRA	90.07 $\pm$ 0.31	17.53 $\pm$ 1.22	72.53 $\pm$ 1.10	89.47 $\pm$ 0.83	58.00 $\pm$ 0.69	31.47 $\pm$ 0.92	92.07 $\pm$ 0.42	17.93 $\pm$ 3.60
<b>GPT-2</b>	Head-Only	85.67 $\pm$ 1.01	6.73 $\pm$ 0.81	78.93 $\pm$ 1.72	82.07 $\pm$ 0.50	47.20 $\pm$ 0.92	34.87 $\pm$ 0.76	85.20 $\pm$ 1.60	10.07 $\pm$ 0.81
	LoRA	91.87 $\pm$ 0.90	7.87 $\pm$ 0.42	84.00 $\pm$ 0.69	90.93 $\pm$ 1.33	54.60 $\pm$ 2.11	36.33 $\pm$ 2.02	92.40 $\pm$ 0.53	7.87 $\pm$ 0.81
<b>Gemma-2B</b>	Head-Only	92.13 $\pm$ 0.58	7.73 $\pm$ 0.64	84.40 $\pm$ 1.22	88.73 $\pm$ 0.46	46.80 $\pm$ 0.20	41.93 $\pm$ 0.50	94.87 $\pm$ 0.23	7.80 $\pm$ 1.06
	LoRA	94.40 $\pm$ 0.53	10.80 $\pm$ 0.53	83.60 $\pm$ 0.69	96.07 $\pm$ 0.31	40.60 $\pm$ 1.73	55.47 $\pm$ 1.70	97.13 $\pm$ 0.50	9.40 $\pm$ 1.51
<b>LLaMA-3.2-1B</b>	Head-Only	93.33 $\pm$ 0.81	7.13 $\pm$ 0.81	86.20 $\pm$ 1.59	86.73 $\pm$ 0.12	47.60 $\pm$ 0.53	39.13 $\pm$ 0.64	94.53 $\pm$ 0.31	10.27 $\pm$ 0.81
	LoRA	93.33 $\pm$ 0.23	8.87 $\pm$ 0.42	84.47 $\pm$ 0.42	95.93 $\pm$ 0.61	43.93 $\pm$ 1.79	52.00 $\pm$ 1.22	97.40 $\pm$ 0.92	8.73 $\pm$ 1.10
<b>Mistral-7B</b>	Head-Only	93.40 $\pm$ 0.53	6.47 $\pm$ 0.50	86.93 $\pm$ 0.95	91.00 $\pm$ 0.53	43.93 $\pm$ 0.58	47.07 $\pm$ 1.01	96.53 $\pm$ 1.55	7.60 $\pm$ 0.72
	LoRA	89.40 $\pm$ 6.05	12.40 $\pm$ 7.50	77.00 $\pm$ 13.53	81.80 $\pm$ 2.12	58.33 $\pm$ 3.75	23.47 $\pm$ 1.75	97.33 $\pm$ 0.61	7.07 $\pm$ 1.68
<b>TinyLLaMA</b>	Head-Only	93.60 $\pm$ 0.92	5.53 $\pm$ 0.31	88.07 $\pm$ 0.81	73.33 $\pm$ 0.12	45.60 $\pm$ 0.20	27.73 $\pm$ 0.31	89.73 $\pm$ 1.86	13.33 $\pm$ 3.01
	LoRA	93.13 $\pm$ 1.03	10.87 $\pm$ 0.42	82.27 $\pm$ 1.40	93.80 $\pm$ 0.20	49.00 $\pm$ 1.59	44.80 $\pm$ 1.64	95.33 $\pm$ 1.22	13.13 $\pm$ 1.63

## E EXPERIMENTAL DETAILS

**Computer-Vision.** For all the experiments, we used a learning rate of  $1e-03$  to train the LoRA and the head-only finetuning, and the training was done using AdaM optimizer (Kingma & Ba, 2014).




Figure 14: Attack Success Rate of the A2T Attack on DistilBERT model when applied to AG News, Banking77, IMDb dataset with different attack budget (number of words changed).

Table 4: Original accuracy (Orig Acc), attack success rate (ASR, %) and adversarial accuracy (Adv Acc, %) across datasets with A2T attack. Attack budget: max number of words changed is 3.

	Dataset	Metric	IMDb			SST-2			Yelp Polarity		
			Orig Acc	ASR	Adv Acc	Orig Acc	ASR	Adv Acc	Orig Acc	ASR	Adv Acc
<b>BERT</b>	Head-Only	86.67 $\pm$ 0.92	8.61 $\pm$ 0.62	79.20 $\pm$ 0.40	83.40 $\pm$ 0.53	28.13 $\pm$ 0.74	59.93 $\pm$ 0.23	87.00 $\pm$ 1.06	11.04 $\pm$ 1.13	77.40 $\pm$ 1.93	
	LoRA	92.93 $\pm$ 1.01	14.85 $\pm$ 0.06	79.13 $\pm$ 0.81	92.13 $\pm$ 0.23	22.86 $\pm$ 0.62	71.07 $\pm$ 0.46	93.20 $\pm$ 0.72	11.52 $\pm$ 0.70	82.47 $\pm$ 1.14	
<b>DistilBERT</b>	Head-Only	87.67 $\pm$ 0.31	6.77 $\pm$ 0.78	81.73 $\pm$ 0.58	82.13 $\pm$ 0.31	28.33 $\pm$ 0.40	58.73 $\pm$ 0.23	86.40 $\pm$ 1.74	11.26 $\pm$ 0.74	76.60 $\pm$ 1.04	
	LoRA	92.87 $\pm$ 0.31	14.86 $\pm$ 1.55	79.07 $\pm$ 1.70	89.47 $\pm$ 0.83	27.80 $\pm$ 1.71	64.60 $\pm$ 1.91	93.33 $\pm$ 1.03	17.50 $\pm$ 2.00	76.93 $\pm$ 2.05	
<b>GPT-2</b>	Head-Only	89.27 $\pm$ 0.95	10.99 $\pm$ 2.04	79.47 $\pm$ 2.39	82.07 $\pm$ 0.50	30.06 $\pm$ 0.42	57.40 $\pm$ 0.40	85.80 $\pm$ 1.91	12.29 $\pm$ 2.30	75.27 $\pm$ 3.06	
	LoRA	93.93 $\pm$ 0.64	13.35 $\pm$ 2.32	81.27 $\pm$ 2.52	90.93 $\pm$ 1.33	22.00 $\pm$ 0.77	70.93 $\pm$ 1.62	92.73 $\pm$ 0.50	8.85 $\pm$ 1.37	84.53 $\pm$ 1.62	

All the tasks were trained for 5 epochs; all of which yielded stable convergence. All the experiments were run 3 times with different seeds to reduce the effect of randomness, and we report the average and the corresponding standard deviations. For the LoRA model, we use the adaptation for all the weights within the self-attention mechanism.

After the convergence of each finetuning strategy, and as explained in Section 6, we used the FGSM and PGD attack with different budget attacks. For the PGD, we set the number of iterations to 10, which is in line with the literature. For the attack budget, we used different attack budgets to illustrate the effect of the attack. Specifically we use a range varying from 4/255 to 32/255. We note that the classical budget for CIFAR-10 and CIFAR-100 is considered to be 8/255, while for the ImageNet it's set to be 4/255.

Finally, for the DeLoRA, we set the  $\lambda$  parameter to 15, following the original work which suggested a value between 10 and 20.

### Natural-Language-Processing.

For both head-only and LoRA fine-tuning, we take the pretrained checkpoint of the model and further finetune it for 5 epochs using the AdamW (Loshchilov & Hutter, 2019) optimizer with a learning rate of  $2e-4$  and batch size of 16. We use the HuggingFace Datasets<sup>1</sup> to get the dataset for the experiment. For the Yelp Polarity dataset, we subsample to 3000 train instances, 1000 validation instances, and 1000 test instances to speed up the training process. For IMDb and SST-2 datasets, we use the full dataset and the original train, validation, and test splits provided by Huggingface.

After the fine-tuning phase, we subsample 500 instances from the test split and run both the TextFooler and A2T attack. The subsampling process is seeded to have a more statistically sound evaluation of the adversarial robustness of head-only and LoRA finetuning. For the TextFooler attack, we allow at most 10 word substitutions per example. Candidate substitutions are drawn from counter-fitted word embeddings (with pre-computed cosine similarity matrix), with up to 50 candidates per word, and are filtered using a sentence-level similarity threshold of 0.7. During model inference for the attack, we use a batch size of 16. For the A2T attack, we allow for the number

<sup>1</sup><https://huggingface.co/docs/datasets/en/index>

1296 of word changes across the set  $\{1, 3, 5, 7\}$  and a query budget of at most 50 forward passes per  
1297 example. Up to 20 candidate substitution words are allowed for each word in a sentence, and the  
1298 substitution words are further narrowed with a word-level similarity of 0.8. Finally, a sentence level  
1299 similarity of 0.9 is used to filter the final sentence. The model inference during the attack is run with  
1300 a batch size of 32.

1301 All experiments are repeated 3 times with different seeds, reporting the average performance to-  
1302 gether with standard deviations.  
1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349