

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DON’T FORGET THE NONLINEARITY: UNLOCKING ACTIVATION FUNCTIONS IN EFFICIENT FINE-TUNING

Anonymous authors

Paper under double-blind review

## ABSTRACT

Existing parameter-efficient fine-tuning (PEFT) methods primarily adapt weight matrices while keeping activation functions fixed. We introduce **NoRA**, the first PEFT framework that directly adapts nonlinear activation functions in pretrained transformer-based models. NoRA replaces fixed activations with learnable rational functions and applies structured low-rank updates to numerator and denominator coefficients, with a group-wise design that localizes adaptation and improves stability at minimal cost. On vision transformers trained on CIFAR-10 and CIFAR-100, NoRA matches or exceeds full fine-tuning while updating only 0.4% of parameters (0.02M), achieving accuracy gains of +0.17% and +0.27%. When combined with LoRA (**NoRA++**), it outperforms LoRA and DoRA under matched training budgets by adding fewer trainable parameters. On LLaMA3-8B instruction tuning, NoRA++ consistently improves generation quality, yielding average MMLU gains of +0.3%–0.8%, including +1.6% on STEM (Alpaca) and +1.3% on OpenOrca. We further show that NoRA constrains adaptation to a low-dimensional functional subspace, implicitly regularizing update magnitude and direction. These results establish activation-space tuning as a complementary and highly parameter-efficient alternative to weight-based PEFT, positioning activation functions as first-class objects for model adaptation.

## 1 INTRODUCTION

Recent advances in deep learning have demonstrated the remarkable power of large-scale pretrained models across domains such as vision, language, and multimodal learning Abnar et al. (2021). However, deploying these models in downstream tasks often requires task-specific adaptation, posing significant challenges in terms of computational efficiency and parameter overhead Jiang et al. (2024); Lyu & Yin (2024). Full fine-tuning of all model parameters is not only costly but also prone to overfitting and catastrophic forgetting, especially when labeled data is limited or hardware resources are constrained Krizhevsky et al. (2009).

To address these issues, parameter-efficient fine-tuning (PEFT) Houlsby et al. (2019) techniques have emerged as a promising solution. Among them, Low-Rank Adaptation (LoRA) Hu et al. (2021) has gained significant attention by introducing trainable low-rank perturbations to frozen weight matrices, achieving strong performance with only a small fraction of trainable parameters. However, while these methods are effective for updating weight matrices, they largely overlook the potential of adapting non-linear components, such as activation functions. Existing PEFT approaches typically treat activation functions as fixed, immutable components, despite their crucial role in capturing task-specific inductive biases (e.g., smoothness, stability) Shi et al. (2024). This neglect of the adaptability of activation functions marks a critical gap in current PEFT strategies. Activations play a vital role in transforming input data at each layer of a neural network Sharma et al. (2017), and their adaptation is key to fine-tuning the model’s performance for specific tasks. This shift from focusing solely on weights to also considering the adaptation of activations represents a fundamental rethinking of the PEFT paradigm, particularly in models like KANs Liu et al. (2025), where the activation functions themselves are learnable and dynamic.

In this work, we investigate fine-tuning strategies that target activation functions, using learnable rational functions as a flexible and expressive alternative to fixed nonlinearities. Unlike traditional architectures that rely on fixed nonlinearities, which remain static during training, our approach

leverages the fact that many widely-used activation functions such as ReLU Glorot et al. (2011), GELU Hendrycks & Gimpel (2016), and Swish Ramachandran et al. (2017), can be closely approximated or even exactly represented using rational functions. Models equipped with learnable rational activations replace these fixed nonlinearities with parameterized rational functions, allowing the nonlinear transformations to be expressed as:

$$\phi(x) = \frac{P(x)}{Q(x)} = \frac{\sum_{i=0}^m a_i x^i}{\sum_{j=0}^n b_j x^j}, \quad (1)$$

where  $\{a_i\}$  and  $\{b_j\}$  are learnable coefficients. This insight implies that our method is theoretically applicable to any network by replacing fixed activations with their rational counterparts. However, adapting rational activations presents unique challenges: small perturbations in the denominator  $Q(x)$  can lead to large functional changes or instability.

To overcome this challenge, we propose the **Nonlinear Rational Adapter (NoRA)**, the first parameter-efficient fine-tuning framework explicitly designed for the activation function components in model. NoRA first replaces the fixed activation functions with learnable rational functions, then introduces low-rank perturbations to both the numerator  $P(x)$  and denominator  $Q(x)$  coefficients, allowing task-specific adaptation while preserving the algebraic structure of rational transformations. By constraining updates to a structured low-dimensional subspace Nie et al. (2020), NoRA ensures smoothness, stability, and bounded functional deviation during fine-tuning—properties critical for the safe adaptation of rational activations.

**Our contributions are summarized as follows:**

- **A new paradigm for activation-centric PEFT:** We introduce NoRA, the first fine-tuning framework that directly targets the adaptation of activation functions. This shifts the focus of PEFT from weight matrices to the nonlinear components of neural networks.
- **Structured low-rank adaptation of rational functions:** NoRA perturbs both numerator and denominator coefficients in a theoretically grounded manner, preserving functional stability while enabling flexible task-specific adaptation.
- **Practical compatibility with rational activations:** NoRA complements existing rational function activations by providing a parameter-efficient adaptation mechanism. It operates without architectural changes, making it readily applicable across models that use rational approximations of standard nonlinearities.

By shifting the focus of PEFT from weight adaptation to activation-level adaptation, our work opens new directions for enhancing expressiveness and adaptability in modern neural architectures.

## 2 RELATED WORK

### 2.1 LOW RANK ADAPTATION (LoRA)

Low-Rank Adaptation (LoRA) Hu et al. (2021) is a technique designed to efficiently fine-tune large pre-trained language models by reducing the number of trainable parameters. Instead of updating the entire weight matrix  $W_0$  during training, LoRA introduces two low-rank matrices,  $A$  and  $B$ , such that:

$$W = W_0 + \Delta W = W_0 + BA \quad (2)$$

Here,  $W_0$  represents the original weight matrix,  $\Delta W$  denotes the weight update,  $B \in \mathbb{R}^{d \times r}$ , and  $A \in \mathbb{R}^{r \times k}$ , where  $r \ll \min(d, k)$  is the rank of the decomposition. This approach leverages the observation that the updates to the weights during model adaptation often have a low intrinsic rank, allowing for a significant reduction in the number of trainable parameters without compromising model performance.

During the forward pass, the output is computed as:

$$y = Wx = (W_0 + BA)x = W_0x + BAx \quad (3)$$

108 In this formulation,  $W_0$  remains fixed, and only the matrices  $A$  and  $B$  are updated during training.  
 109 This strategy not only reduces computational and memory requirements but also mitigates issues  
 110 such as catastrophic forgetting by preserving the original model parameters.  
 111

112 **2.2 LEARNABLE ACTIVATION FUNCTIONS**  
 113

114 Activation functions are critical to the expressivity and inductive biases of neural networks.  
 115 While standard architectures rely on fixed nonlinearities such as ReLU Glorot et al. (2011),  
 116 GELU Hendrycks & Gimpel (2016), and Swish Ramachandran et al. (2017), recent work has  
 117 explored learnable activation functions that adapt their shape during training. Early parametric  
 118 forms include PReLU He et al. (2015a), which learns a slope parameter for negative activations,  
 119 and APL Agostinelli et al. (2015), which models activations as a piecewise-linear combination of  
 120 hinge functions. Later developments use spline-based and kernel-based approximations for higher  
 121 flexibility.  
 122

123 A particularly powerful family of learnable activations is based on rational functions. Due to their  
 124 universal approximation property Baker Jr & Gammel (1961), rational functions can represent a  
 125 wide range of continuous functions more compactly than polynomials. A rational activation is  
 126 typically expressed as:  
 127

$$\phi(x) = \frac{P(x)}{Q(x)} = \frac{\sum_{i=0}^m a_i x^i}{\sum_{j=0}^n b_j x^j}, \quad (4)$$

128 where  $\{a_i\}$  and  $\{b_j\}$  are learnable coefficients of the numerator and denominator, respectively. This  
 129 formulation allows activation functions to dynamically adjust their nonlinearity during training.  
 130

131 **3 NONLINEAR RATIONAL ADAPTER (NoRA)**  
 132

133 In this work, we propose the **Nonlinear Rational Adapter (NoRA)**, a novel parameter-efficient  
 134 fine-tuning method that adapts pretrained models by modifying their nonlinear activation functions  
 135 with learnable rational functions.  
 136

137 Traditional activation functions, such as ReLU Glorot et al. (2011) and GELU Hendrycks & Gimpel  
 138 (2016), can all be approximately expressed as rational functions Telgarsky (2017) of the form:  
 139

$$\phi(x) = \frac{P(x)}{Q(x)} \quad (5)$$

140 where  $P(x)$  and  $Q(x)$  are polynomials of the form  $P(x) = \sum_{i=0}^m a_i x^i$  and  $Q(x) = \sum_{j=0}^n b_j x^j$ ,  
 141 with learnable coefficients  $\{a_i\}$  and  $\{b_j\}$ . Then this formulation also can be standardized to avoid  
 142 division by zero as:  
 143

$$\phi(x) = \frac{a_0 + a_1 x + a_2 x^2 + \cdots + a_m x^m}{1 + |b_0 + b_1 x + b_2 x^2 + \cdots + b_n x^n|} \quad (6)$$

144 This formulation allows any fixed activation function to be represented as a rational function with  
 145 learnable coefficients, providing more flexibility and expressiveness in modeling complex data trans-  
 146 formations. While using a single shared rational activation function for all neurons limits the model’s  
 147 expressiveness, assigning a unique activation to each neuron is prohibitively expensive. Following  
 148 Group-KAN Yang & Wang (2025), we partition the hidden (channel) dimension of each layer into  $g$   
 149 disjoint groups, fixed across tokens and batches. All neurons in the same group share one learnable  
 150 rational activation function. This static grouping preserves flexibility while keeping the overhead  
 151 linear in  $g$  (i.e.,  $g$  activations per layer) rather than per neuron.  
 152

153 Building on this idea, NoRA first replaces the fixed activation function with a group learnable rational  
 154 function, and then injects structured low-rank perturbations Benaych-Georges & Nadakuditi  
 155 (2011) into the coefficient matrices of these rational functions. Specifically, the perturbations are  
 156 applied to both the numerator  $P$  and the denominator  $Q$  in a grouped fashion, where the coefficients  
 157 are divided into  $g$  groups. Let  $\phi(X)$  denote the original rational function, and  $\phi'(X)$  its perturbed  
 158 counterpart. The resulting updated rational function is given by:  
 159

$$\phi'(X) = \frac{(P_g + \mathcal{L}_g(\Delta P))(X)}{(Q_g + \mathcal{L}_g(\Delta Q))(X)} \quad (7)$$

162 Here,  $P_g$  and  $Q_g$  represent the original polynomial numerator and denominator of the rational activation  
 163 function, respectively. The perturbation terms  $\Delta P$  and  $\Delta Q$  are approximated using a group-  
 164 wise low-rank adaptation function  $\mathcal{L}_g(\cdot)$ , which applies independent low-rank perturbations within  
 165 each group  $g = 1, \dots, G$ .

166 More concretely, the entire set of neurons is partitioned into  $G$  disjoint groups, each containing  $n/G$   
 167 neurons. For each group  $g$ , the adapted rational activation function is defined as  
 168

$$169 \phi'_g(X) = \frac{(P_g + A_g^P B_g^P)(X)}{(Q_g + A_g^Q B_g^Q)(X)} \quad (8)$$

$$170$$

$$171$$

172 where the perturbations  $\Delta P_g = A_g^P B_g^P$  and  $\Delta Q_g = A_g^Q B_g^Q$  are expressed as low-rank matrix  
 173 products with

$$174 A_g^{(\cdot)} \in \mathbb{R}^{d \times r}, \quad B_g^{(\cdot)} \in \mathbb{R}^{r \times 1},$$

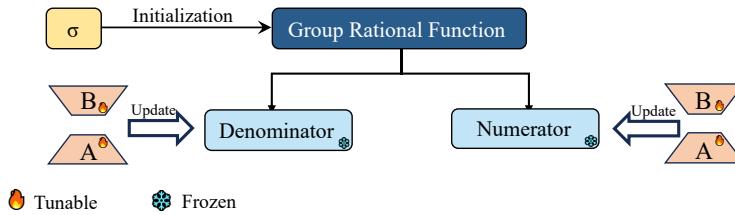
$$175$$

176 where  $(\cdot) \in \{P, Q\}$ ,  $d$  is the degree of the polynomial, and  $r$  is the rank of the approximation.  
 177

178 All neurons within the same group  $g$  share the adapted activation function  $\phi'_g$ , enabling parameter-  
 179 efficient and localized functional adaptation. Increasing the number of groups  $G$  enhances the gran-  
 180 uularity of adaptation, allowing more flexible modeling of complex activation patterns while keeping  
 181 the parameter increase minimal due to the low-rank structure.  
 182

183 Furthermore, to ensure stable training and smooth fine-tuning from a pretrained baseline, all low-  
 184 rank matrices are initialized similarly to LoRA Hu et al. (2021):  $A_g$  is initialized with a small Gaus-  
 185 sian noise (e.g.,  $\mathcal{N}(0, 0.02)$ ) and  $B_g$  is initialized as zeros. This initialization guarantees that the  
 186 adapted rational activation functions  $\phi'_g(X)$  are equivalent to the original functions  $P_g(X)/Q_g(X)$   
 187 at the start of training. The reason why tuning activation matters and why we choose rational func-  
 188 tion are shown in Appendix A, B.

189 An overview of the NoRA framework is illustrated in Figure 1.



190 Figure 1: Overview of the NoRA framework. NoRA replaces fixed activation functions with group  
 191 rational functions and introduces structured low-rank perturbations to both the numerator and de-  
 192 nominator coefficients.  
 193

## 200 4 EXPERIMENT

201 In this section, we evaluate the performance of NoRA across both image classification and language  
 202 model tasks. Specifically, we apply NoRA to the ViT-Tiny model for CIFAR-10 and CIFAR-100  
 203 classification, and to the LLaMA3-8B model for instruction tuning.  
 204

### 205 4.1 COMPARISON WITH PARAMETER-EFFICIENT FINE-TUNING METHODS IN IMAGE 206 CLASSIFICATION

#### 207 4.1.1 EXPERIMENT SETUP.

208 We conduct experiments using the ViT-Tiny model pretrained on ImageNet-1K Deng et al. (2009)  
 209 For adaptation, we explore two PEFT configurations: (1) **NoRA**, where we only replace the GELU in  
 210 FFN with the rational activation function and use group-wise low-rank perturbations while keeping  
 211 all other weights frozen; and (2) **NoRA++**, a hybrid variant that combines NoRA with standard  
 212 LoRA applied to the MLP layers in attention layers. In NoRA++, both the activation functions  
 213

216 Table 1: Comparison with parameter-efficient fine-tuning methods on CIFAR-10 and CIFAR-100.  
217

218 <b>Method</b>	219 <b>Trainable Params (M)</b>	220 <b>CIFAR-10 Acc. (%)</b>	221 <b>CIFAR-100 Acc. (%)</b>
220 Full tuning	221 5.54 (100%)	222 90.71	223 77.19
221 VPT	222 0.39 (7.0%)	223 89.62 ( <b>-1.09</b> )	224 75.43 ( <b>-1.76</b> )
222 Adapter	223 0.48 (8.7%)	224 89.93 ( <b>-0.78</b> )	225 75.88 ( <b>-1.31</b> )
223 LoRA	224 0.33 (6.0%)	225 91.05 ( <b>+0.34</b> )	226 77.68 ( <b>+0.49</b> )
224 QLoRA	225 0.33 (6.0%)	226 90.45 ( <b>-0.26</b> )	227 77.01 ( <b>-0.18</b> )
225 DoRA	226 0.34 (6.1%)	227 91.13 ( <b>+0.42</b> )	228 77.71 ( <b>+0.52</b> )
226 <b>NoRA (ours)</b>	227 0.02 (0.4%)	228 90.88 ( <b>+0.17</b> )	229 77.46 ( <b>+0.27</b> )
227 <b>NoRA++ (ours)</b>	228 0.35 (6.2%)	229 <b>91.24</b> ( <b>+0.53</b> )	230 <b>77.76</b> ( <b>+0.57</b> )

228 and select linear weights are jointly adapted, offering a more expressive yet still parameter-efficient  
229 fine-tuning scheme. For both NoRA and NoRA++, the low-rank perturbation rank is set to  $r = 2$ .  
230 In both settings, the classification head is also trained. Evaluation is performed on CIFAR-10 and  
231 CIFAR-100 Krizhevsky et al. (2009), two widely used image classification benchmarks. CIFAR-  
232 10 includes 10 object classes, while CIFAR-100 contains 100 fine-grained classes grouped into 20  
233 superclasses. Each dataset provides 6000 or 600 samples per class, with image resolution  $32 \times 32$ .  
234 We resize images to  $224 \times 224$  and use a patch size of 16 Dosovitskiy et al. (2020) during training.  
235 The specific hyperparameter settings can be referred to in Appendix C.  
236

#### 237 4.1.2 BASELINE METHODS.

238 To provide a comprehensive evaluation, we compare NoRA with several representative parameter-  
239 efficient fine-tuning (PEFT) methods. These include **Full Fine-Tuning**, which updates all model  
240 parameters and serves as an upper-bound reference; **VPT** Jia et al. (2022), which prepends learnable  
241 visual prompt tokens to the input sequence while keeping the backbone frozen; **Adapter** Chen et al.  
242 (2022), which inserts lightweight bottleneck modules between transformer blocks and updates only  
243 these modules during training; **LoRA** Hu et al. (2021), which introduces trainable low-rank matrices  
244 into attention layers while freezing the original weights; **QLoRA** Dettmers et al. (2023), which ex-  
245 tends LoRA to 4-bit quantized models for memory-efficient adaptation; and **DoRA** Liu et al. (2024),  
246 which decomposes pre-trained weights into magnitude and direction components to better approxi-  
247 mate the behavior of full fine-tuning. All methods are implemented on the same ViT-Tiny backbone  
248 for fair comparison, with the classification head remaining trainable. Detailed hyperparameter set-  
249 tings are provided in Appendix C.1.  
250

#### 251 4.1.3 RESULT ANALYSIS

252 As shown in Table 1, NoRA, while tuning only 0.02M parameters (0.4%), achieves 90.88% accuracy  
253 on CIFAR-10 and 77.46% on CIFAR-100, outperforming full fine-tuning by +0.17% and +0.27%,  
254 respectively. This highlights the surprising effectiveness of adaptively tuning activation functions  
255 alone, without modifying any backbone weights. In contrast, other PEFT baselines such as LoRA  
256 and DoRA also slightly surpass full fine-tuning but require over 6% of the model parameters to be  
257 updated—more than 15× as many as NoRA. Meanwhile, Adapter, QLoRA, and VPT lag behind in  
258 both accuracy and efficiency, underscoring the importance of adaptation position and mechanism.  
259 To explore composability, we further introduce a hybrid variant, **NoRA++**, which applies NoRA to  
260 the activation functions and LoRA to the attention and MLP layers. This integration yields the best  
261 accuracy on both datasets—91.24% on CIFAR-10 and 77.76% on CIFAR-100—while still using  
262 fewer trainable parameters than full fine-tuning. These results confirm that NoRA offers an excellent  
263 trade-off between accuracy and efficiency, and its compatibility with weight-based methods like  
264 LoRA enables scalable and flexible adaptation strategies.  
265

#### 266 4.2 SCALABILITY WITH GROUP EXPANSION.

##### 267 4.2.1 EXPERIMENT SETUP.

268 To further investigate the capacity and scalability of NoRA, we systematically increase the number  
269 of groups  $g$  by powers of two, setting  $g = 8, 16, 32, 64$ , while keeping the rank fixed at  $r = 3$  and all

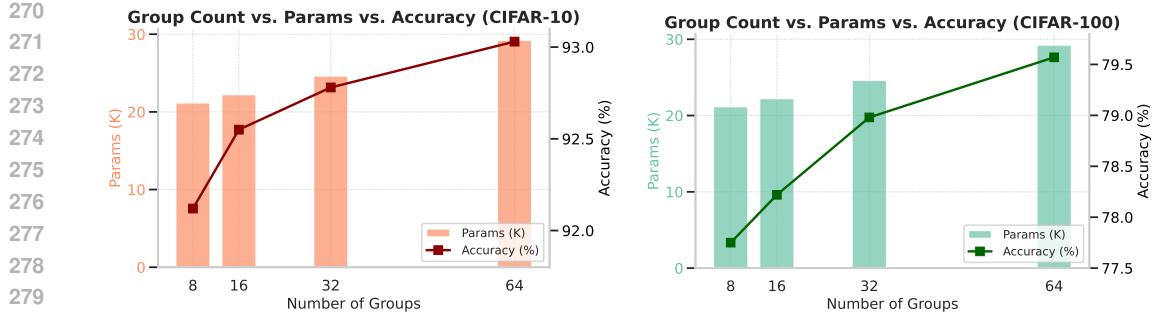


Figure 2: The relationship curve between the number of groups and the number of parameters (without classification head) and the classification accuracy of CIFAR-10/100 during fine-tuning.

other training settings unchanged. This allows us to study how finer group partitioning influences the expressiveness and adaptability of the rational activation functions. To initialize the parameters for these finer group divisions efficiently, we adopt a simple replication strategy, wherein the original group-wise rational coefficients are duplicated along the channel dimension to match the increased number of groups. This approach enables higher-resolution activation adaptation without modifying the model architecture, thus providing a practical and scalable way to enhance NoRA’s functional flexibility.

#### 4.2.2 RESULT ANALYSIS.

This straightforward strategy yields consistent performance improvements, as illustrated in Figure 2. As the number of groups  $g$  increases from 8 to 64, NoRA achieves progressively higher accuracy on both CIFAR-10 (from approximately 92.2% to 93.1%) and CIFAR-100 (from approximately 77.6% to 79.6%). Meanwhile, the number of trainable parameters increases only moderately (from approximately 13K to 28K), reflecting the controlled growth enabled by the low-rank group-wise design. These results empirically confirm that increasing group resolution facilitates more localized and specialized nonlinear modeling, allowing NoRA to better capture task-specific activation dynamics. The dual-axis plots in Figure 2 clearly illustrate this trade-off: performance scales almost linearly with group count, while parameter cost grows sub-linearly, underscoring NoRA’s efficiency. Notably, these gains are achieved without tuning additional hyperparameters or introducing significant training overhead, further highlighting the practicality and scalability of the proposed method. This behavior supports the intuition that composing structured local approximations can effectively approximate global nonlinear functions within the activation space.

### 4.3 INSTRUCT-TUNING IN LARGE LANGUAGE MODEL

#### 4.3.1 EXPERIMENT SETUP.

To assess the compatibility and enhanced effectiveness of our method in joint instruction-tuning settings, we introduce NoRA++, a hybrid adaptation framework that combines the proposed activation-centric fine-tuning with conventional weight-space tuning. Specifically, we integrate NoRA with LoRA on the LLaMA3-8B model by replacing the fixed activation functions in the MLP blocks of each Transformer layer with group-wise rational functions, and then applying structured low-rank adaptation via NoRA. This modification enables activation-space tuning while maintaining the parameter efficiency of LoRA’s weight-space updates, resulting in a complementary and synergistic tuning mechanism. We evaluate NoRA++ on a suite of five diverse instruction datasets—*Alpaca*, *MathInstruct*, *OpenOrca*, *ShareGPT-Hyper*, and *UltraChat*—encompassing tasks from open-ended dialogue and reasoning to summarization. Under identical training budgets, NoRA++ consistently surpasses standard LoRA across all datasets, achieving significant improvements in output quality and generalization. To quantitatively assess its reasoning ability, we report results on the MMLU Hendrycks et al. (2020) benchmark under a 5-shot evaluation setup. NoRA++ yields consistent gains of +0.3% to +0.8% in average accuracy over LoRA, demonstrating that activation-centric

adaptation provides meaningful benefits even when layered atop established PEFT methods. Full implementation and hyperparameter details are provided in Appendix C .

### 4.3.2 RESULT ANALYSIS.

As shown in Table 2, the combination of NoRA and LoRA yields consistent and measurable performance gains across most instruction-tuning datasets and MMLU categories. For example, on Alpaca and MathInstruct, NoRA++ improves the average accuracy by +0.8% and +0.5%, respectively, with notable gains such as +1.6% in STEM (Alpaca) and +2.3% in STEM (MathInstruct). Even on OpenOrca, which already benefits from strong LoRA tuning, the addition of NoRA leads to further improvements in key areas like STEM (+1.3%) and Social Sciences (+1.2%), resulting in a net average increase. While a few categories exhibit minor regressions—for instance, -0.7% in Humanities on OpenOrca and -0.4% in the “Other” category of ShareGPT-Hyper—the overall trend remains positive. These results suggest that NoRA effectively introduces complementary local nonlinearity at the activation level, enhancing model expressiveness without disrupting the core low-rank structure imposed by LoRA. This synergy between activation-level and weight-level tuning highlights NoRA’s general applicability and its potential as a plug-and-play enhancement module for a wide range of parameter-efficient fine-tuning (PEFT) methods in large-scale language models.

Table 2: MMLU-test accuracy (%) after instruction tuning LLaMA3-8B on five datasets. Each cell shows LoRA result followed by NoRA+LoRA result and delta. Green indicates improvement and red indicating decline.

Tuning Dataset	STEM	Humanities	Social Sciences	Other	Average
Alpaca	60.2 → <b>61.8</b> (+1.6)	65.5 → <b>65.0</b> (-0.5)	63.3 → <b>64.5</b> (+1.2)	62.0 → <b>63.0</b> (+1.0)	62.8 → <b>63.6</b> (+0.8)
MathInstruct	53.7 → <b>56.0</b> (+2.3)	75.5 → <b>74.9</b> (-0.6)	58.3 → <b>59.2</b> (+0.9)	70.7 → <b>70.4</b> (-0.3)	63.9 → <b>64.4</b> (+0.5)
OpenOrca	54.2 → <b>55.5</b> (+1.3)	74.6 → <b>73.9</b> (-0.7)	58.1 → <b>59.3</b> (+1.2)	71.0 → <b>70.4</b> (-0.6)	63.9 → <b>64.2</b> (+0.3)
ShareGPT-Hyper	58.9 → <b>59.2</b> (+0.3)	66.3 → <b>67.0</b> (+0.7)	60.7 → <b>61.0</b> (+0.3)	59.4 → <b>59.0</b> (-0.4)	61.3 → <b>61.6</b> (+0.3)
UltraChat	57.2 → <b>57.9</b> (+0.7)	64.8 → <b>65.3</b> (+0.5)	60.2 → <b>59.8</b> (-0.4)	59.7 → <b>60.6</b> (+0.9)	60.5 → <b>60.9</b> (+0.4)

### 4.4 ABLATION STUDY AND ANALYSIS

In this section we present four experiments to assess rational activations, low-rank perturbations, selective coefficient tuning, and overall efficiency.

#### 4.4.1 COMPARISON WITH OTHER LEARNABLE ACTIVATIONS.

To validate the necessity of employing rational functions as activation mechanisms, we perform an ablation study comparing NoRA with several common learnable activation functions. Specifically, we replace the GELU Hendrycks & Gimpel (2016) activations in the MLP layers of the pretrained ViT-Tiny model with alternative nonlinearities, including **PReLU** He et al. (2015b), **ARelu** Chen et al. (2020), and **ELU** Clevert et al. (2016). We then fine-tune only the activation function parameters and the classification head on CIFAR-100, while keeping all other model weights frozen. As reported in Table 3, these general-purpose activation functions yield only marginal improvements, with top-1 accuracy consistently lagging behind our proposed rational activation-based approach. This suggests that simple substitution of activation functions fails to provide sufficient task-specific adaptability or structural compatibility within the frozen transformer architecture. In contrast, our method enables structured, localized, and task-adaptive modulation of the activation landscape through low-rank perturbations of rational functions, yielding superior representational refinement under stringent parameter constraints.

Table 3: Ablation on learnable activation functions.

Name	Accuracy (%)
PReLU	53.21
ARelu	54.17
ELU	52.84
<b>NoRA (Ours)</b>	<b>77.46</b>

#### 4.4.2 DIFFERENT TUNING METHODS FOR LEARNABLE RATIONAL ACTIVATIONS.

To evaluate the effectiveness of NoRA, we compare it with several fine-tuning strategies that all keep the backbone frozen and only update the classification head along with the learnable rational

378 functions on CIFAR-100. **Rational fine-tuning** directly updates the coefficients of the rational activation  
 379 functions. **Zero-initialized fine-tuning** introduces a learnable matrix initialized to zero and  
 380 added to the rational coefficients. **GELU-initialized fine-tuning** initializes the rational functions to  
 381 approximate GELU before training. **Const-tuning (a<sub>0</sub>,b<sub>0</sub> only)** restricts adaptation to the constant  
 382 terms of the numerator and denominator, i.e., updating only  $a_0$  and  $b_0$ , which controls global off-  
 383 set/scale but leaves the nonlinear shape fixed. As shown in Table 4, all three baselines underperform,  
 384 suggesting that coefficient-only tuning lacks sufficient expressiveness. In contrast, NoRA achieves  
 385 the highest accuracy by introducing a learnable low-rank shift in the activation space, offering more  
 386 flexible and effective adaptation.

#### 388 4.4.3 IMPACT OF SELECTIVE PERTURBATION ON RATIONAL COEFFICIENTS.

389 To further investigate the importance of jointly  
 390 adapting both components of the rational activation,  
 391 we conduct an ablation study in which low-rank per-  
 392 turbations are selectively applied to either the nu-  
 393 merator or the denominator coefficients, while keep-  
 394 ing the other component fixed. As presented in Ta-  
 395 ble 5, perturbing only one side leads to a notice-  
 396 able performance degradation on both CIFAR-10  
 397 and CIFAR-100, compared to the setting where both  
 398 components are jointly adapted. This result suggests  
 399 that the effectiveness of NoRA stems not from mod-  
 400 ulating a single part of the activation function, but  
 401 from the synergistic interaction between the nu-  
 402 merator and denominator. Specifically, the numerator controls the functional shape of the activation,  
 403 while the denominator governs numerical stability and saturation behavior. Their co-adaptation  
 404 introduces a richer and more flexible activation landscape, which is critical for the improved gen-  
 405 eralization performance observed. These findings underscore the necessity of NoRA’s co-perturbation  
 406 design and highlight a fundamental architectural distinction from existing parameter-efficient fine-  
 407 tuning methods.

Table 4: Ablation on different fine-tuning strategies.

Method	Accuracy (%)
Rational tuning	76.56
Zero-init tuning	76.44
GELU-init tuning	76.07
Const-tuning (a <sub>0</sub> ,b <sub>0</sub> only)	74.91
<b>NoRA (ours)</b>	<b>77.46</b>

Numerator	Denominator	CIFAR-10 Acc. (%)	CIFAR-100 Acc. (%)
✗	✓	87.40	74.16
✓	✗	86.92	73.58
✓	✓	<b>90.88</b>	<b>77.46</b>

Table 5: Effect of selectively injecting low-rank perturbations into numerator and denominator co-  
 efficients. ✓ indicates perturbed (trainable), ✗ indicates unperturbed (frozen).

#### 418 4.4.4 RESOURCE EFFICIENCY ANALYSIS.

419 We assess the resource efficiency of NoRA and LoRA from three key perspectives: inference-time  
 420 computational cost (FLOPs), latency, and the number of trainable parameters, all evaluated on the  
 421 ViT-Tiny backbone. As shown in Table 6, NoRA achieves over 17 $\times$  reduction in trainable parameter  
 422 count compared to LoRA (2.10K vs. 37.24K), highlighting its extreme parameter efficiency. Despite  
 423 this, the inference latency remains comparable—5.69 ms/sample for NoRA vs. 5.30 ms/sample for  
 424 LoRA—indicating that the additional expressiveness introduced by activation modulation does not  
 425 impose substantial runtime overhead. The slight increase in FLOPs (1.07G vs. 0.91G) is expected,  
 426 as NoRA introduces group-wise rational activation functions, which involve evaluating both poly-  
 427 nomial numerators and denominators across multiple activation groups. This minor computational  
 428 overhead is a direct result of NoRA’s more flexible nonlinear modeling capability, and reflects the  
 429 trade-off between functional expressiveness and cost. Importantly, the trainable parameter counts re-  
 430 ported here exclude the classification head to ensure fair comparison. Overall, these results reinforce  
 431 NoRA’s strength as a lightweight and deployment-friendly PEFT method, with minimal runtime cost  
 and strong potential for integration with other techniques like LoRA or adapters.

Method	Params (K)	FLOPs (G)	Inference Time (ms/sample)
LoRA	37.24	0.91	5.30
<b>NoRA (ours)</b>	2.10	1.07	5.69

Table 6: Comparison of resource efficiency between LoRA and NoRA on ViT-Tiny.

## 5 CONCLUSION AND FUTURE WORK

### 5.1 CONCLUSION

In this work, we proposed the **Nonlinear Rational Adapter (NoRA)**, a novel and general framework for parameter-efficient fine-tuning that shifts the focus of adaptation from the traditional weight-centric view to a new activation-centric perspective. By replacing fixed nonlinearities with task-adaptive rational functions, NoRA enables flexible and expressive modulation of pretrained models through compact, structured perturbations applied directly in the activation space. This is achieved via learnable low-rank parameterizations of both the numerator and denominator of rational functions, allowing for stable, efficient, and interpretable task adaptation while preserving the overall model architecture. The proposed group-wise formulation further enhances scalability by localizing adaptation across independently modulated activation units, thereby reducing parameter redundancy and improving representational flexibility. Extensive experiments on image classification (CIFAR-10/100) and instruction tuning of large language models (LLaMA3-8B) demonstrate that NoRA significantly outperforms full fine-tuning, despite updating only 0.4% of the total model parameters. Moreover, our extended variant, NoRA++, which integrates NoRA with LoRA for joint adaptation of activations and weights, achieves even stronger performance, consistently surpassing both DoRA and LoRA across vision and language tasks while maintaining superior parameter efficiency. Collectively, these results validate activation-centric adaptation as a powerful and underexplored dimension in the fine-tuning landscape, offering a complementary perspective to weight-based PEFT approaches and paving the way for more modular, robust, and efficient model customization strategies within large-scale pretraining paradigms.

### 5.2 FUTURE WORK

We plan to extend NoRA beyond classification and instruction tuning to more complex architectures, including generative diffusion models, graph neural networks, and vision and language models, where structured nonlinearity may be particularly beneficial. We will study functional classes beyond rationals (for example spline based, Fourier inspired, or attention conditioned activations) to clarify expressiveness and efficiency tradeoffs. We will develop adaptive group wise strategies that automatically select the number of groups  $g$  and the subspace rank  $r$  during training, which enables dynamic control of capacity and cost. We will also integrate NoRA with complementary PEFT methods such as prompt tuning, adapters, and LoRA variants across different network levels to achieve compositional and task aware adaptation. Finally, we aim to scale NoRA to trillion parameter language models and to evaluate it on long context, multi hop reasoning, and multi task benchmarks. Because NoRA operates in activation space and is modular with respect to weight centric updates, it can be inserted into existing pipelines and is suitable for deployment under resource or privacy constraints, including edge inference, online learning, and user personalization, where adapting lightweight nonlinear components enables efficient and stable continual learning on dynamic or streaming data.

486 ETHICS STATEMENT  
487488 This work adheres to the ICLR Code of Ethics. Our study does not involve human subjects, sensitive  
489 data, or potentially harmful applications. All experiments were conducted using publicly available  
490 datasets, and there are no conflicts of interest or ethical concerns associated with this research.  
491492 REPRODUCIBILITY STATEMENT  
493494 We provide detailed descriptions of datasets, experimental settings, and model hyperparameters in  
495 the main text and appendix. All algorithms and theoretical claims are fully documented, and code  
496 to reproduce the results will be made publicly available upon publication.  
497498 REFERENCES  
499500 Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, and Hanie Sedghi. Exploring the limits of  
501 large scale pre-training, 2021. URL <https://arxiv.org/abs/2110.02095>.  
502503 Forest Agostinelli, Matthew Hoffman, Peter Sadowski, and Pierre Baldi. Learning activation func-  
504 tions to improve deep neural networks, 2015. URL <https://arxiv.org/abs/1412.6830>.  
505506 George A Baker Jr and John L Gammel. The padé approximant. *Journal of Mathematical Analysis*  
507 and *Applications*, 2(1):21–30, 1961.  
508509 Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low  
510 rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.  
511512 Dengsheng Chen, Jun Li, and Kai Xu. Arelu: Attention-based rectified linear unit, 2020. URL  
513 <https://arxiv.org/abs/2006.13858>.  
514515 Zhe Chen, Yuchen Duan, Wenhui Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision  
516 transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022.  
517518 Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network  
519 learning by exponential linear units (elus), 2016. URL <https://arxiv.org/abs/1511.07289>.  
520521 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-  
522 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
523 pp. 248–255. Ieee, 2009.  
524525 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning  
526 of quantized llms. *Advances in neural information processing systems*, 36:10088–10115, 2023.  
527528 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
529 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An  
530 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint  
arXiv:2010.11929*, 2020.  
531532 Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In  
533 *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp.  
534 315–323. JMLR Workshop and Conference Proceedings, 2011.  
535536 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing  
537 human-level performance on imagenet classification. In *Proceedings of the IEEE international  
538 conference on computer vision*, pp. 1026–1034, 2015a.  
539540 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing  
541 human-level performance on imagenet classification, 2015b. URL <https://arxiv.org/abs/1502.01852>.  
542

540 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint*  
 541 *arXiv:1606.08415*, 2016.

542

543 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and  
 544 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint*  
 545 *arXiv:2009.03300*, 2020.

546

547 Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, An-  
 548 drea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp.  
 549 In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.

550

551 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 552 and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021. URL <https://arxiv.org/abs/2106.09685>.

553

554 Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and  
 555 Ser-Nam Lim. Visual prompt tuning. In *European conference on computer vision*, pp. 709–727.  
 556 Springer, 2022.

557

558 Ziyue Jiang, Bo Yin, and Boyun Lu. Precise apple detection and localization in orchards using  
 559 yolov5 for robotic harvesting systems. In *2024 2nd International Conference on Algorithm, Image  
 Processing and Machine Vision (AIPMV)*, pp. 262–265. IEEE, 2024.

560

561 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.  
 562 2009.

563

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

567

568 Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić,  
 569 Thomas Y. Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks, 2025. URL <https://arxiv.org/abs/2404.19756>.

570

571 Yilin Lyu and Bo Yin. A discussion of migration of common neural network regularization methods  
 572 on snns. In *Ninth International Symposium on Advances in Electrical, Electronics, and Computer  
 573 Engineering (ISAEECE 2024)*, volume 13291, pp. 1355–1361. SPIE, 2024.

574

575 Feiping Nie, Wei Chang, Zhanxuan Hu, and Xuelong Li. Robust subspace clustering with low-rank  
 576 structure constraint. *IEEE Transactions on Knowledge and Data Engineering*, 34(3):1404–1415,  
 577 2020.

578

579 Prajit Ramachandran, Barret Zoph, and Quoc V Le. Swish: a self-gated activation function. *arXiv  
 580 preprint arXiv:1710.05941*, 7(1):5, 2017.

581

582 Sagar Sharma, Simone Sharma, and Anidhya Athaiya. Activation functions in neural networks.  
 583 *Towards Data Sci*, 6(12):310–316, 2017.

584

585 Yiming Shi, Jiwei Wei, Yujia Wu, Ran Ran, Chengwei Sun, Shiyuan He, and Yang Yang. Loldu:  
 586 Low-rank adaptation via lower-diag-upper decomposition for parameter-efficient fine-tuning.  
 587 *arXiv preprint arXiv:2410.13618*, 2024.

588

589 Matus Telgarsky. Neural networks and rational functions. In *International Conference on Machine  
 590 Learning*, pp. 3387–3393. PMLR, 2017.

591

592 Xingyi Yang and Xinchao Wang. Kolmogorov-arnold transformer. In *The Thirteenth International  
 593 Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=BCeock53nt>.

594

595 Longteng Zhang, Lin Zhang, Shaohuai Shi, Xiaowen Chu, and Bo Li. Lora-fa: Memory-efficient  
 596 low-rank adaptation for large language models fine-tuning. *arXiv preprint arXiv:2308.03303*,  
 597 2023.

594 **A WHY TUNING ACTIVATION MATTERS?**  
595596 **Setup and notation.** Consider a depth- $L$  network with frozen weight matrices  $W_1, \dots, W_L$  and  
597 elementwise, possibly group-shared, activations  $\phi_\ell(\cdot; \theta_\ell)$ :

598 
$$h_0(x) = x, \quad z_\ell = W_\ell h_{\ell-1}(x), \quad h_\ell(x) = \phi_\ell(z_\ell; \theta_\ell), \quad F(x) = h_L(x). \quad (9)$$
  
599

600 Here  $\theta_\ell \in \mathbb{R}^{p_\ell}$  are *activation parameters*. We analyze why adapting  $\{\theta_\ell\}$  (while freezing  $\{W_\ell\}$ ) is  
601 impactful for expressivity, stability, optimization, and generalization.602 **A.1 FUNCTIONAL DIRECTIONS UNLOCKED BY ACTIVATION PARAMETERS**  
603604 Let  $\Delta\theta = \{\Delta\theta_\ell\}_{\ell=1}^L$  be a small update. By the chain rule,  
605

606 
$$F(x; \theta + \Delta\theta) - F(x; \theta) = \sum_{\ell=1}^L \underbrace{\frac{\partial F}{\partial h_\ell}(x)}_{\Pi_{j=\ell+1}^L J_j(x)} \cdot \underbrace{\frac{\partial h_\ell}{\partial \theta_\ell}(x)}_{D_\ell(x) \Delta\theta_\ell} + \mathcal{O}(\|\Delta\theta\|^2). \quad (10)$$
  
607  
608  
609

610 We denote

611 
$$J_j(x) := \text{Diag}(\phi'_j(z_j; \theta_j)) W_j, \quad (11)$$
  
612 and

613 
$$D_\ell(x) \in \mathbb{R}^{d_\ell \times p_\ell}, \quad [D_\ell(x)]_{i,:} = \left[ \frac{\partial \phi_\ell(z_{\ell,i}; \theta_\ell)}{\partial \theta_\ell} \right]^\top. \quad (12)$$
  
614

615 Thus the *first-order functional change* lies in the span

616 
$$\Delta F(\cdot) \in \text{span} \left\{ (\Pi_{j=\ell+1}^L J_j(\cdot)) D_\ell(\cdot) e_k : 1 \leq \ell \leq L, 1 \leq k \leq p_\ell \right\}. \quad (13)$$
  
617

618 Its intrinsic dimension is at most  $\sum_\ell p_\ell$  (or  $\sum_\ell G_\ell p_\ell$  if parameters are shared across  $G_\ell$  groups).  
619 Therefore, activation tuning provides a low-dimensional yet function-space set of directions unavailable if activations are fixed.  
620621 **A.2 STABILITY AND LIPSCHITZ CONTROL**  
622623 Denote the Lipschitz seminorm of  $\phi_\ell(\cdot; \theta_\ell)$  as

624 
$$\text{Lip}(\phi_\ell; \theta_\ell) := \sup_z |\phi'_\ell(z; \theta_\ell)|. \quad (14)$$
  
625

626 For each block  $h_\ell = \phi_\ell \circ W_\ell$ ,

627 
$$\text{Lip}(h_\ell; \theta_\ell) \leq \text{Lip}(\phi_\ell; \theta_\ell) \|W_\ell\|_2. \quad (15)$$
  
628

629 Hence the network Lipschitz constant satisfies

630 
$$\text{Lip}(F; \theta) \leq \prod_{\ell=1}^L \text{Lip}(\phi_\ell; \theta_\ell) \|W_\ell\|_2. \quad (16)$$
  
631  
632  
633

634 **Theorem A.1** (Network-level deviation under activation updates). *Let  $F'$  denote the network after  
635 changing  $\theta \mapsto \theta + \Delta\theta$ . Then for any  $x$ ,*

636 
$$\|F'(x) - F(x)\|_2 \leq \sum_{\ell=1}^L \left( \prod_{j>\ell} \text{Lip}(h_j; \theta_j) \right) \|\Delta\phi_\ell(z_\ell)\|_2, \quad (17)$$
  
637  
638  
639

640 where

641 
$$\Delta\phi_\ell(z) := \phi_\ell(z; \theta_\ell + \Delta\theta_\ell) - \phi_\ell(z; \theta_\ell). \quad (18)$$
  
642

643 Moreover,

644 
$$\|\Delta\phi_\ell(z_\ell)\|_2 \leq \sup_z \|D_\ell(z)\|_{2 \rightarrow 2} \|\Delta\theta_\ell\|_2 \sqrt{d_\ell} + \mathcal{O}(\|\Delta\theta_\ell\|_2^2). \quad (19)$$
  
645

646 *Consequences.* (i) By directly controlling  $\text{Lip}(\phi_\ell; \theta_\ell)$  one can regularize the global Lipschitz constant,  
647 which enters standard generalization and robustness bounds. (ii) The deviation bound depends linearly (first order) on activation parameters through  $D_\ell$ , enabling fine-grained, stable modulation even with frozen weights.

648 A.3 GRADIENT FLOW, JACOBIANS, AND EFFECTIVE RANK  
649650 For a single block  $h = \phi(Wx; \theta)$ , the input Jacobian is  
651

652 
$$J_x = \frac{\partial h}{\partial x} = \text{Diag}(\phi'(Wx; \theta)) W. \quad (20)$$
  
653

654 Activation parameters change  $J_x$  via  $\phi'$ , thus altering (i) sensitivity to inputs, (ii) conditioning of  
655 the layer, and (iii) gradient flow to preceding layers. If  $\phi$  is gated (e.g., with a slope/threshold  
656 parameter), let  $\mathcal{A}(\theta) := \{i : \phi'(z_i; \theta) > 0\}$  be the active set. Then  
657

658 
$$\text{Rank}(J_x) \leq |\mathcal{A}(\theta)| \leq d. \quad (21)$$
  
659

660 Hence, tuning  $\theta$  modulates the effective rank and spectrum of the linearized map  $\Pi_\ell J_\ell$ , improving  
661 gradient propagation in deep stacks.  
662663 A.4 OPTIMIZATION GEOMETRY AND THE NTK PERSPECTIVE  
664665 Let  $\vartheta := (\{W_\ell\}, \{\theta_\ell\})$  and write the neural tangent kernel (NTK)  $K_\vartheta(x, x') =$   
666  $\langle \nabla_\vartheta F(x), \nabla_\vartheta F(x') \rangle$ . Decompose  
667

668 
$$K_\vartheta(x, x') = K_{WW}(x, x') + K_{\theta\theta}(x, x') + 2K_{W\theta}(x, x'). \quad (22)$$
  
669

670 Using the first-order expansion,  
671

672 
$$\nabla_{\theta_\ell} F(x) = \left( \Pi_{j > \ell} J_j(x) \right) D_\ell(x), \quad (23)$$
  
673

674 so  $K_{\theta\theta}$  spans activation-feature directions  $\{D_\ell\}$  propagated to the output.  
675676 **Proposition A.1** (Complementarity at initialization). *Suppose pre-activations  $z_\ell$  are centered and  
677 whitened within groups, and the initial  $W_\ell$  are independent of  $\theta_\ell$  (with zero-mean entries). Then*  
678

679 
$$\mathbb{E} K_{W\theta}(x, x') = 0, \quad \mathbb{E} K_\vartheta(x, x') = \mathbb{E} K_{WW}(x, x') + \mathbb{E} K_{\theta\theta}(x, x'). \quad (24)$$
  
680

681 A.5 CURVATURE CONTROL VIA  $\phi'$  AND  $\phi''$   
682683 Let  $\mathcal{L}$  be a twice-differentiable loss and denote  $J_x = \partial h / \partial x$ . For block  $h = \phi(Wx; \theta)$ ,  
684

685 
$$\frac{\partial^2 h}{\partial x^2} = \text{Diag}(\phi''(Wx; \theta)) [Wx] [Wx]^\top \odot I. \quad (25)$$
  
686

687 Hence the input Hessian of the loss obeys  
688

689 
$$\nabla_x^2 \mathcal{L} = J_x^\top \nabla_h^2 \mathcal{L} J_x + \sum_i \frac{\partial \mathcal{L}}{\partial h_i} \frac{\partial^2 h_i}{\partial x^2}. \quad (26)$$
  
690

691 Tuning  $\theta$  therefore changes both the Gauss–Newton part (through  $\phi'$  in  $J_x$ ) and the residual curvature  
692 term (through  $\phi''$ ), reshaping the local optimization landscape without touching  $W$ .  
693694 A.6 DATA-ALIGNED GRADIENTS AND SATURATION AVOIDANCE  
695696 Let the gradient w.r.t. a preceding weight matrix  $W_\ell$  be  
697

698 
$$\nabla_{W_\ell} \mathcal{L} = (\delta_\ell \odot \phi'_\ell(z_\ell; \theta_\ell)) h_{\ell-1}(x)^\top, \quad (27)$$
  
699

700 where  $\delta_\ell$  is the backpropagated error. Then  
701

702 
$$\mathbb{E} \|\nabla_{W_\ell} \mathcal{L}\|_F^2 = \mathbb{E} [\|\delta_\ell\|_2^2 \|h_{\ell-1}\|_2^2] \cdot \mathbb{E} \overline{\phi'_\ell(z_\ell; \theta_\ell)^2}, \quad (28)$$
  
703

704 where the overline denotes the average across units. Choosing  $\theta_\ell$  to *maximize derivative mass* under  
705 the moments of  $z_\ell$  keeps most units responsive (not saturated), improving gradient signal-to-noise  
706 without changing  $W_\ell$ .  
707

702 A.7 GENERALIZATION VIA CAPACITY CONTROL  
703704 Define the hypothesis class  $\mathcal{F}$  with frozen  $\{W_\ell\}$  and activation parameters bounded by  $\|\theta_\ell\| \leq \rho_\ell$   
705 and local Lipschitz budgets  $\text{Lip}(\phi_\ell; \theta_\ell) \leq \lambda_\ell$ . Using standard Lipschitz-based complexity bounds,  
706 for inputs  $\|x\| \leq R$  and  $N$  samples,

707 
$$\mathfrak{R}_N(\mathcal{F}) \lesssim \frac{R}{\sqrt{N}} \cdot \left( \sum_{\ell=1}^L \left( \prod_{j>\ell} \lambda_j \|W_j\|_2 \right) \cdot c_\ell(\rho_\ell) \|W_\ell\|_2 \right), \quad (29)$$
  
708

709 where  $c_\ell(\rho_\ell)$  is a polynomial-in- $\rho_\ell$  constant induced by the chosen parametrization of  $\phi_\ell$ .  
710711 Takeaways.  
712713

- **Functional expressivity at low cost.** Activation parameters open a low-dimensional  
714 *function-space* of directions (Eq. 13).
- **Stable, controllable modulation.** Activation Lipschitz and curvature  $(\phi', \phi'')$  yield explicit  
715 network-level deviation and robustness control (Eqs. 16, 17, 26).
- **Better conditioning and gradient flow.** By changing gates/slopes, activation tuning  
716 controls Jacobian rank/spectrum and mitigates saturation (Eqs. 20, 21, 28).
- **Complementary optimization geometry.** Activation-parameter gradients contribute an  
717 NTK component largely orthogonal (in expectation) to weight-only directions (Eqs. 22,  
718 23, 24).

719 These results justify *activation tuning* as a principled and effective axis for adapting pretrained  
720 models, independently of the particular parameterization chosen for  $\phi_\ell$ .  
721722 B WHY RATIONAL FUNCTION?  
723724 Rational functions offer a compact and flexible parametrization that can *uniformly approximate* the  
725 activation functions used in modern networks—both smooth (e.g., tanh, sigmoid, GELU/erf-based,  
726 SiLU/Swish) and non-smooth (e.g., ReLU, Leaky/ReLU)—on any bounded pre-activation domain.  
727 We record the key facts concisely.728 **Definition.** A degree- $(m, n)$  rational function is  
729

730 
$$r_{m,n}(x) = \frac{P_m(x)}{Q_n(x)} = \frac{\sum_{i=0}^m a_i x^i}{\sum_{j=0}^n b_j x^j}, \quad (30)$$
  
731

732 and we assume the domain  $K \subset \mathbb{R}$  is compact with a pole-free margin  
733

734 
$$\inf_{x \in K} |Q_n(x)| \geq \gamma > 0, \quad (31)$$
  
735

736 which is standard for stable approximation on  $K$ .  
737738 **Density on compact sets.** Because polynomials are dense in  $C(K)$  (Stone–Weierstrass) and poly-  
739 nomials are a special case of rationals (take  $Q_n \equiv 1$ ), rationals are also dense:  
740

741 
$$\forall f \in C(K), \forall \varepsilon > 0, \exists m, n, \exists r_{m,n} \text{ s.t. } \sup_{x \in K} |f(x) - r_{m,n}(x)| < \varepsilon. \quad (32)$$
  
742

743 Thus any continuous activation used in practice admits uniform rational approximations on bounded  
744 pre-activation ranges.  
745746 **Fast rates for smooth activations.** When  $f$  is real-analytic on a neighborhood of  $K$  (typical for  
747 sigmoid/tanh/erf-like activations), best rational approximants achieve geometric convergence:  
748

749 
$$\exists C > 0, \rho > 1 : \inf_{\deg(r) \leq N} \sup_{x \in K} |f(x) - r(x)| \leq C \rho^{-N}, \quad (33)$$
  
750

751 where  $N = m + n$  is total degree. This is substantially faster than the algebraic rates of many  
752 polynomial schemes.  
753

756    **Near-root-exponential rates for kinks.** For non-smooth activations with a finite number of kinks  
 757    (e.g., ReLU,  $|x|$ ), rational approximation still excels:

$$759 \quad \exists C, c > 0 : \inf_{\deg(r) \leq N} \sup_{x \in K} |f(x) - r(x)| \leq C \exp(-c\sqrt{N}), \quad (34)$$

760    whereas the best polynomial error decays only algebraically in  $N$ . Hence even piecewise-linear  
 761    activations can be approximated to high accuracy with modest rational degree.

763    **Practical consequences.**

- 765    • **Coverage.** Eqs. equation 32–equation 34 ensure a single rational family can approximate  
 766    most activations used in deep learning on bounded pre-activation sets.
- 767    • **Efficiency.** The fast rates in equation 33–equation 34 imply that low degrees suffice in  
 768    practice, keeping parameters and compute small.
- 769    • **Stability.** The pole-free margin equation 31 ensures bounded slopes/curvatures on  $K$ , mak-  
 770    ing training numerically stable while retaining expressive shape control via the coefficients.

772    **C EXPERIMENT DETAILS**

774    The tables 7 and table 8 below show the hyperparameter during tuning the models.

776    Table 7: Hyperparameter settings for training ViT-Tiny on CIFAR-100.

778 <b>Hyperparameter</b>	779 <b>Value</b>
780    Input resolution	224 <sup>2</sup>
781    Epochs	50
782    Batch size	256
783    Learning rate	$1 \times 10^{-3}$
784    Learning rate decay	Cosine
785    Optimizer	AdamW
786    Weight decay	0.05
787    AMP	True

789    Table 8: Hyperparameter settings for instruct tuning.

791 <b>Hyperparameter</b>	792 <b>Value</b>
793    Cutoff length	1024
794    Flash attention	auto
795    Max new tokens	512
796    Max samples	1000
797    Per-device eval batch size	2
798    Preprocessing workers	16
799    Quantization method	bnn
800    Stage	SFT
801    Temperature	0.95
802    Template	default
803    Top-p	0.7
804    Trust remote code	True

805    **C.1 IMPLEMENTATION DETAILS OF BASELINE METHODS**

807    To ensure fair comparison across all parameter-efficient fine-tuning (PEFT) methods, we adopt a  
 808    unified experimental setup based on the ViT-Tiny backbone pretrained on ImageNet-1K. All meth-  
 809    ods fine-tune the classification head, and images are resized to  $224 \times 224$  using a patch size of 16.  
 Below we detail the specific hyperparameter configurations for each baseline:

- **VPT** Jia et al. (2022): We prepend 10 learnable prompt tokens of dimension 192 to the input sequence. Only the prompts and classification head are updated. The learning rate for the prompts is  $5 \times 10^{-3}$ , and the training schedule matches that of full fine-tuning.
- **Adapter** Houlsby et al. (2019): Adapter modules with a bottleneck dimension of 48 are inserted between each transformer block. Only adapter parameters and the classification head are updated. Learning rate is set to  $1 \times 10^{-3}$ .
- **LoRA** Zhang et al. (2023): LoRA modules with rank  $r = 8$  are inserted into the query and value projections of each attention layer. Alpha is set to 16, and dropout is disabled. Only LoRA parameters and the classification head are updated.
- **QLoRA** Dettmers et al. (2023): The backbone is quantized to 4-bit precision using NF4 format. LoRA is applied with the same configuration as above. We use gradient checkpointing and double quantization as described in the original paper.
- **DoRA** Liu et al. (2024): All linear weights are decomposed into magnitude and direction, with both components trainable. Initialization follows the frozen pre-trained weights. The learning rate is  $5 \times 10^{-4}$ , consistent with the original DoRA setup.

## D OTHER EXPERIMENT RESULTS

### D.1 OTHER ABLATION STUDIES

**Rank Setting.** We evaluate the performance of NoRA under varying amounts of tunable parameters by adjusting the rank  $r$  in the low-rank updates of the rational function coefficients. Specifically, we use the ViT-Tiny model pretrained on ImageNet-1K and fix the rational function structure like GR-KAN as  $(m = 5, n = 4)$ . To investigate the trade-off between expressivity and parameter efficiency, we experiment with  $r \in \{1, 2, 3, 4\}$  while keeping other training configurations unchanged. As shown in Figure 3, increasing the rank leads to improved performance up to  $r = 3$ , beyond which gains saturate. The setting  $r = 3$  achieves the best accuracy on CIFAR-100, suggesting that it strikes a good balance between capacity and efficiency. Notably, when calculating the number of parameter in this task we ignore the classification head.

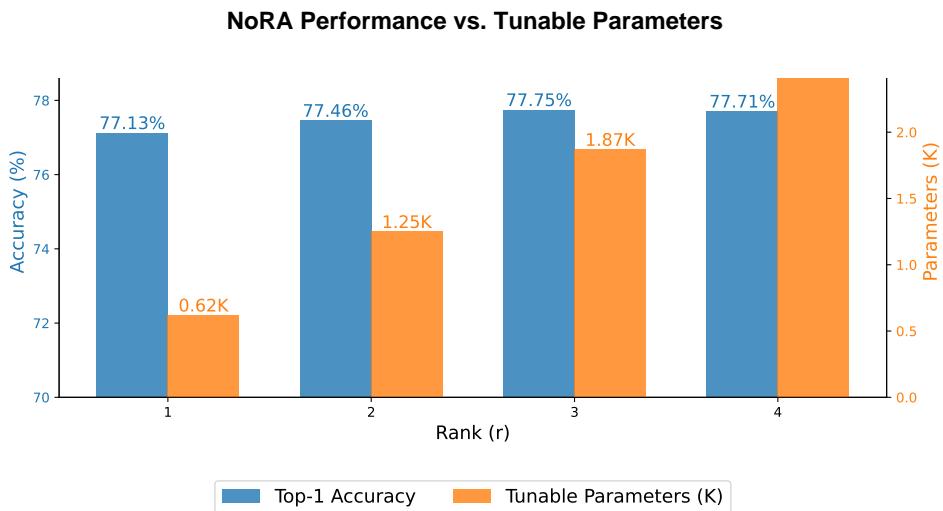
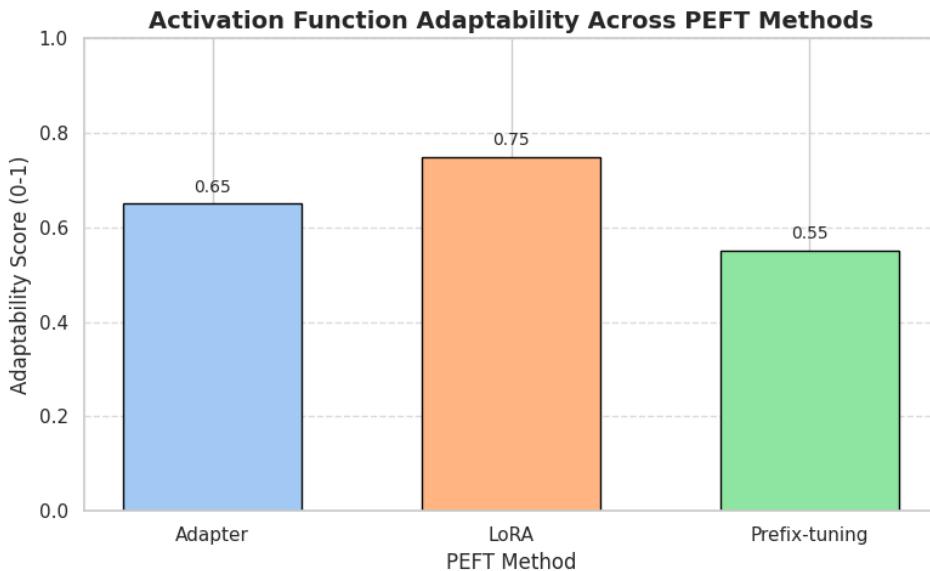


Figure 3: Accuracy-Parameter Trade-off in NoRA on CIFAR-100 with Varied Low-Rank Updates on ImageNet-1K Pretrained ViT-Tiny.

864 D.2 ADAPTABILITY OF ACTIVATION FUNCTIONS WITH DIFFERENT PEFT METHODS  
865

866 Before and after fine-tuning the model, if the activation distribution of the same batch of samples  
867 significantly changes across layers, it indicates that the method provides the activation function with  
868 greater *plasticity*. Conversely, if the activation distribution remains relatively stable, the adaptability  
869 is weaker. We typically use distribution distance as the core metric, normalizing it to the interval  
870  $[0, 1]$ , and define the **Adaptability Score** in this context. As shown in Figure 4, mainstream methods  
871 have not yet effectively explored the adaptability of activation functions.



892 Figure 4: Comparative analysis of Different PEFT Methods in terms of the adaptability of activation  
893 functions

894  
895 D.3 ANALYSIS OF CONVERGENCE RATE

896 As illustrated in Figure 5, we plot the training curves of NoRA and full fine-tuning across training  
897 epochs. NoRA converges significantly faster, stabilizing around epoch 20, whereas full fine-tuning  
898 requires nearly 45 epochs to reach convergence. This demonstrates NoRA’s capacity to leverage pre-  
899 trained knowledge more efficiently, leading to faster and more stable adaptation. Likewise, although  
900 our method does not outperform full fine-tuning on the training set, it achieves better generalization  
901 on the test set. This indicates that our approach imposes a beneficial inductive bias, likely mitigating  
902 overfitting and preserving useful priors from the pretrained model.

903  
904 D.4 RESULT VISUALIZATION

905 **t-SNE Visualization.** To further understand the representational effect of NoRA, we perform a  
906 t-SNE visualization on the learned feature embeddings for a selected subset of 10 diverse CIFAR-  
907 100 classes. As shown in Figure 6, the embeddings obtained through full fine-tuning exhibit well-  
908 separated clusters, indicating task-specific adaptation with high discriminative power. Interestingly,  
909 NoRA achieves a similarly clear cluster structure despite tuning only 0.02% of parameters, suggest-  
910 ing that it successfully reshapes the learned representation space within a low-dimensional subspace.

911 **Grad-CAM Visualization.** To intuitively demonstrate how adjusting nonlinear activations influ-  
912 ences transferability, we visualize the final block of the ViT-Tiny model before and after fine-tuning  
913 using Grad-CAM on selected samples from CIFAR-100. As shown in Figure 7, the fine-tuned  
914 model exhibits stronger focus on the main objects within the images, indicating enhanced feature  
915 localization. This suggests that fine-tuning the activation functions effectively improves the model’s  
916 performance on downstream tasks by enabling better transfer and reuse of relevant features.

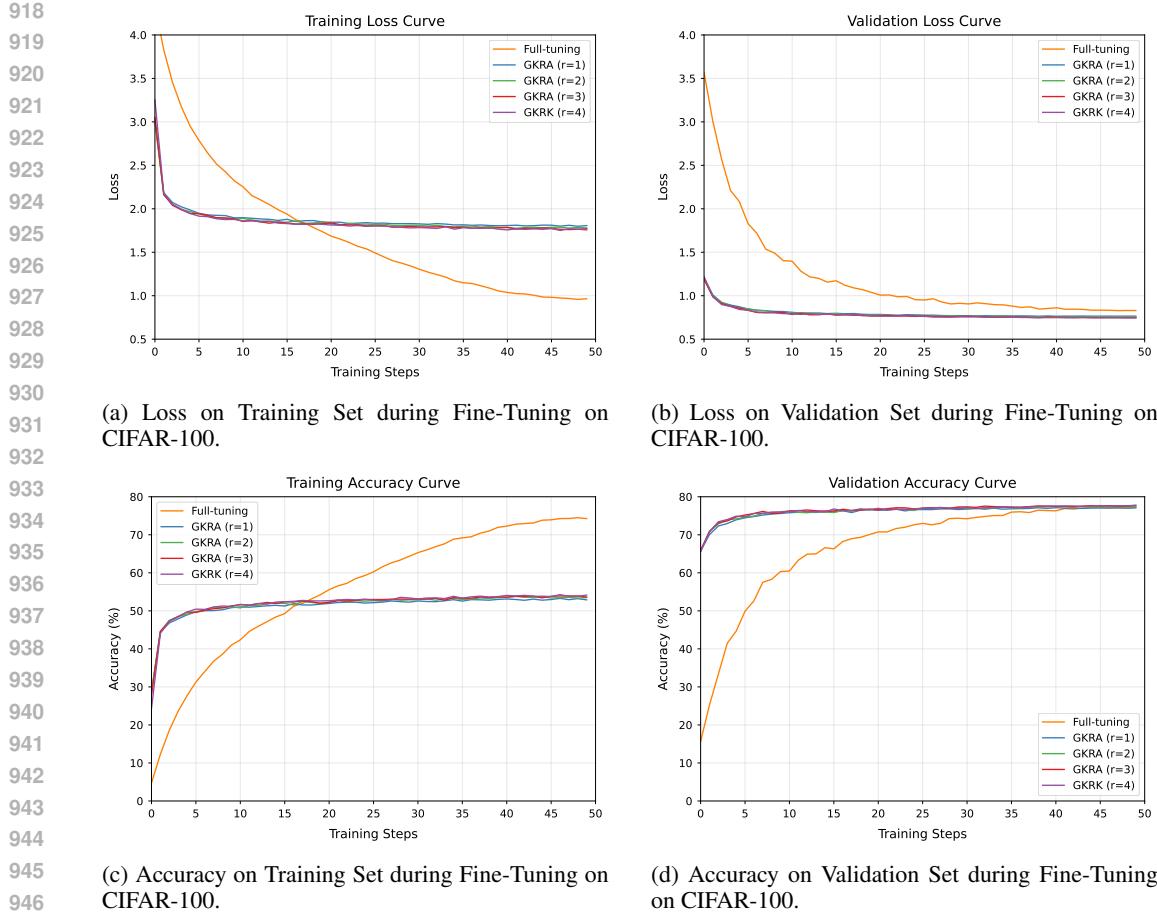


Figure 5: Comparison of Convergence Efficiency and Accuracy between NoRA and Full Fine-Tuning across four experimental settings.

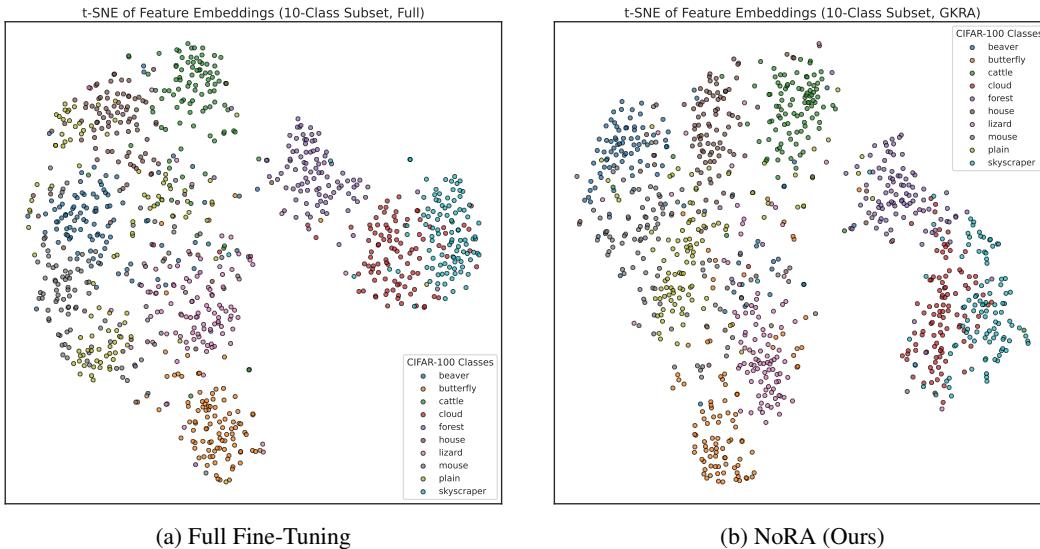


Figure 6: t-SNE visualization of feature embeddings on a 10-class subset of CIFAR-100. (a) Full fine-tuning produces well-separated clusters. (b) NoRA achieves comparably structured representations with 0.02% parameter updates, illustrating its capacity to retain discriminative geometry while preserving pretrained inductive priors.

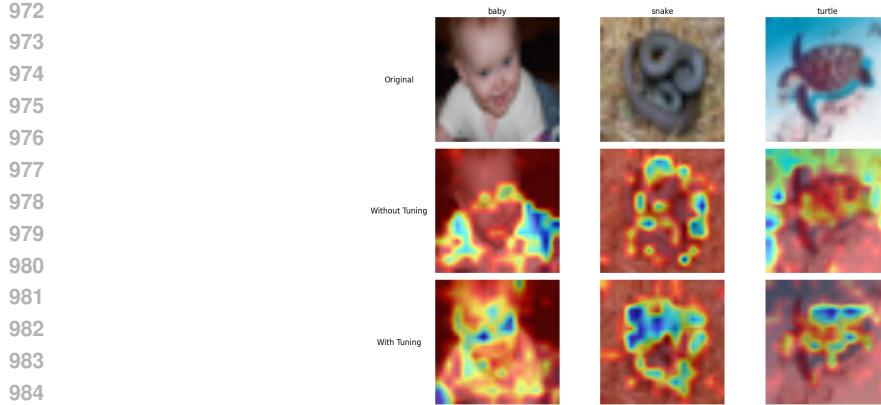


Figure 7: Grad-CAM comparison chart before and after NoRA fine-tuning

## E CODE AVAILABLE

Our code is available in [https://anonymous.4open.science/r/NoRA\\_1](https://anonymous.4open.science/r/NoRA_1).

## F THE USE OF LARGE LANGUAGE MODELS (LLMs)

Parts of this manuscript were linguistically polished with the assistance of a large language model (LLM), specifically ChatGPT (GPT-5). The model was only used for improving grammar, phrasing, and clarity. All research ideas, experimental designs, data collection, analyses, and conclusions are solely the work of the authors.

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