

# 000 001 002 003 004 005 THE CHOSEN FEW: SPARSE ADAPTATION FOR LARGE 006 MODELS 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035

## ABSTRACT

036 Parameter-Efficient Fine-Tuning (PEFT) methods have become essential for adapt-  
037 ing large pretrained models to downstream tasks, with Low-Rank Adaptation  
038 (LoRA) emerging as one of the most widely adopted solutions. However, there  
039 remain several key limitations in current LoRA-based PEFT methods: (1) the low-  
040 rank feature space in LoRA is rigid, reducing its capacity for dynamic adaptation;  
041 (2) the restricted dimensionality, coupled with dense and entangled representations,  
042 constrains the model’s capacity to generalize across multiple domains; and (3) the  
043 compression process limits the extent to which model behavior can be understood  
044 from the learned representations, making it difficult to interpret the functional role  
045 of task-relevant features. In this paper, we argue that *sparse adaptation* offers a  
046 principled and more flexible alternative to low-rank adaptation, with the added  
047 benefit of enhancing interpretability. Instead of compressing information into a low-  
048 rank subspace, sparse adaptation focuses on identifying and selectively activating  
049 a small subset of high-dimensional latent features, enabling a more decomposed  
050 and dynamic fine-tuning process. Building on this paradigm, we propose STAN  
051 (**S**parse **a**dap**T**atio**N**), a novel method that actualizes sparse adaptation by integrat-  
052 ing dedicated Sparse Autoencoder (SAE) modules into frozen pretrained models.  
053 STAN learns to encode task-specific adaptations through sparse activations within  
054 the SAEs, thereby using sparse features as the mechanism for dynamic and robust  
055 adaptation. Beyond the flexibility offered by input-dependent sparse combinations,  
056 the large latent space of the SAEs provides scalable capacity for cross-domain  
057 adaptation, while their inherent semantic decomposition structure supports more  
058 interpretable representations. Through extensive experiments, we demonstrate that  
059 STAN outperforms state-of-the-art PEFT baselines across a range of benchmarks,  
060 while uniquely enabling inspection and analysis of the learned sparse activations.  
061 Our findings position sparse adaptation as a promising new direction in PEFT,  
062 advancing both the expressivity and interpretability of model adaptation.  
063  
064

## 065 1 INTRODUCTION

066 Large-scale pretrained models have demonstrated strong generalization across a wide range of  
067 modalities, including language, vision, and multimodal tasks. Representative examples include large  
068 language models (LLMs) (Bai et al., 2023; Brown et al., 2020; Guo et al., 2025), vision-language  
069 models (VLMs) (Liu et al., 2023b), diffusion-based image generators (Ho et al., 2020; Rombach et al.,  
070 2022), and large vision transformers (Dosovitskiy et al., 2020), which have achieved state-of-the-art  
071 zero-shot and few-shot performance in a wide range of downstream tasks (Kojima et al., 2022). These  
072 capabilities are largely attributed to pretraining on massive and heterogeneous datasets, enabling  
073 models to internalize broad statistical regularities and encode diverse semantic knowledge. Despite  
074 their versatility, adapting such large foundation models, often containing billions of parameters, to new  
075 tasks or domains requires substantial computational resources and specialized expertise. Full model  
076 fine-tuning typically requires extensive task-specific supervision, incurs significant computational  
077 cost, and can lead to undesirable side effects such as catastrophic forgetting (Luo et al., 2023)  
078 and memorization of sensitive data (Carlini et al., 2019). These challenges have motivated the  
079 development of parameter-efficient fine-tuning (PEFT) methods (Houlsby et al., 2019; Hu et al.,  
080 2022; Li & Liang, 2021; Xu et al., 2023), which aim to adapt models by modifying only a small  
081 subset of parameters while maintaining performance.  
082  
083

Among PEFT methods (Houlsby et al., 2019; Hu et al., 2022; Li & Liang, 2021; Xu et al., 2023), Low-Rank Adaptation (LoRA) (Hu et al., 2022) has gained widespread adoption due to its simplicity and strong empirical performance. By inserting trainable low-rank matrices into frozen pretrained layers, LoRA enables efficient adaptation with minimal overhead, often achieving results comparable to full model fine-tuning across a range of downstream tasks. However, LoRA’s adaptation mechanism introduces several key limitations: (1) The low-rank feature space inherent in LoRA (Hu et al., 2022) imposes structural rigidity, significantly limiting the model’s capacity for dynamic adaptation to diverse data distributions (Wang & Zhao, 2025). This inflexibility hinders the alignment of adapted features with the specific and evolving demands of heterogeneous tasks and data characteristics. (2) The restricted dimensionality, often coupled with the emergence of dense and entangled representations, substantially constrains the model’s ability to generalize across distinct domains (Zhang et al., 2025). The lack of clear separability in these compressed features diminishes the model’s capacity to learn domain-specific nuances and adapt meaningfully in multi-domain or cross-distribution contexts. (3) The inherent compression in LoRA’s mechanism reduces the interpretability of the learned representations (Nijasure et al., 2025). This opacity makes it difficult to interpret the functional roles of task-relevant features, limiting our ability to identify which components drive adaptation, and thus posing a challenge to understanding model behaviors. These limitations collectively lead to reduced adaptability, insufficient domain generalization, and limited capacity to revise, debug, or exert fine-grained control over the adaptation process. In addition, the limited understanding of model behavior (Chen et al., 2025) falls short of addressing the growing demand for *interpretable adaptation* in large models – a need that has been increasingly emphasized in recent work (Mumuni & Mumuni, 2025; Wang et al., 2025).

To address the gap above, we propose a shift from *low-rank compression* to *sparse adaptation* – a framework grounded in sparse feature learning (Cunningham et al., 2023; Olshausen & Field, 1997). Rather than encoding adaptation into rigid, dense, and polysemantic subspaces, STAN seeks to *dynamically* activate a sparse subset of high-dimensional, disentangled features that more directly reflect task-specific changes in model behavior. Sparse Autoencoders (SAEs) (Cunningham et al., 2023; Makhzani & Frey, 2013) provide a natural mechanism for this goal, enabling models to learn robust representations with sparsity constraints that promote structure separability and semantic decoupling. Furthermore, the broad representation capacity and the sparsity rendered by SAEs can provide the adaptation with more selection combinations, which not only enhances the *dynamics* of fine-tuning but also strengthens the model’s ability to generalize across diverse domains.

We instantiate this framework with STAN (Sparse adapTatioN), a new PEFT method designed to enhance the efficiency and dynamics of adaptation, expand the overall representational capacity, and enable flexible, scalable adaptation with improved interpretability. Rather than projecting into low-rank dense subspaces, STAN integrates lightweight SAE modules into selected layers of a frozen pretrained model. These modules encode adaptation signals using sparsely activated high-dimensional features, offering a richer representational space **with only modest overhead addition**. STAN achieves adaptation efficiency comparable to LoRA while enabling more scale feasible feature space selection and more broad representation capacity over the adaptation process, facilitating the adaptation process to capture the diverse decomposed information in multi-space combination. We validate STAN across a diverse range of tasks and architectures, demonstrating that it achieves competitive performance with strong PEFT baselines. Moreover, besides the reasoning tasks, we extend STAN to diffusion-based generative models (Ho et al., 2020; Rombach et al., 2022), where we show that the learned sparse features support qualitative and quantitative analysis, enabling structured interventions such as multi-style alignment and interpretability with disentangled features. An overview of STAN’s architecture is provided in Figure 1. We summarize our main contributions as follows:

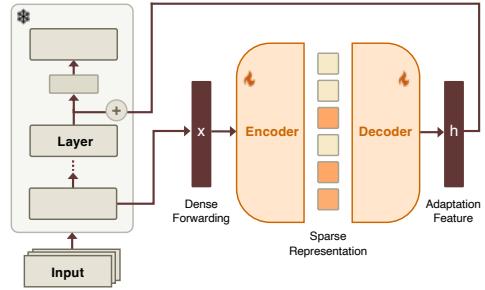


Figure 1: **Overview of STAN pipeline.**  $x$  is the dense forwarding and  $h$  denotes the adaptation feature to the next layer. The layer block is a frozen pretrained model layer. The neuron between the Encoder and Decoder is latent feature, and only the top  $K$  activations (orange) are passed to the decoder, while the rest are masked out.

- 108 • We introduce STAN, a novel PEFT method that instantiates the Sparse Adaptation paradigm  
109 using Sparse Autoencoders, offering a conceptually distinct alternative to low-rank adapta-  
110 tion methods such as LoRA, and enabling a more dynamic and flexible fine-tuning process.
- 111 • We demonstrate the effectiveness of STAN across a range of tasks and model architectures  
112 spanning multiple modalities, showing that it achieves performance superior to state-of-the-  
113 art baselines. Empirical analysis further reveals that STAN retains adaptation stability while  
114 improving the dynamics of the fine-tuning process through the aid of sparsity.
- 115 • We extend STAN to a state-of-the-art diffusion-based generative model (Rombach et al.,  
116 2022) in the context of a style alignment task, showcasing its versatility and further high-  
117 lighting its large representational capacity to capture and isolate diverse styles. We aim to  
118 enhance interpretability by dynamically decomposing dense features into sparse, disentan-  
119 gled components. By understanding the semantics of each component and their interactions,  
120 we hope to be able to reason about the behavior of the entire adaptation mechanism.

## 122 2 RELATED WORK

124 **Parameter-Efficient Fine-Tuning Methods.** Parameter-Efficient Fine-Tuning has emerged as  
125 an effective strategy for adapting large pretrained models to downstream tasks through dedicated  
126 modifications (Houlsby et al., 2019). Rather than updating all model parameters, PEFT focuses  
127 on adjusting a small subset, significantly reducing computational and data requirements. PEFT  
128 methods can be broadly categorized into three classes (Prottasha et al., 2025): prompt tuning, adapter  
129 modules, and low-rank decomposition adaptation methods: (i) prompt tuning (Lester et al., 2021; Liu  
130 et al., 2023b; 2024a; Shen et al., 2024a) steers the model outputs by optimizing either continuous  
131 or discrete prompts while keeping the backbone frozen. These prompts are typically injected into  
132 the input or intermediate layers and act as virtual tokens that guide the model’s predictions. A  
133 prominent example is Prefix-Tuning (Li & Liang, 2021), which prepends trainable vectors into the  
134 attention mechanism, enabling effective task adaptation with minimal parameter updates; (ii) Adapter  
135 modules (Houlsby et al., 2019; He et al., 2021; Hu et al., 2023) introduce additional lightweight  
136 layers within pretrained architectures. These modules are often structured as bottlenecks, comprising  
137 a down-projection, non-linear activation, and an up-projection, allowing the core model to remain  
138 untouched. AdapterHub (Pfeiffer et al., 2020) exemplifies this modularity, supporting plug-and-play  
139 adapters across tasks; (iii) *Low-rank decomposition adaptation methods* (Hu et al., 2022; Liu et al.,  
140 2023c; Tian et al., 2024; Xia et al., 2024) take a different approach by approximating weight updates  
141 using low-rank matrices. LoRA (Hu et al., 2022) is the canonical method, freezes the base model and  
142 injects trainable low-rank matrices into specific layers. Several extensions have since been proposed:  
143 AdaLoRA (Zhang et al., 2023) dynamically allocates parameter budgets based on learned importance  
144 scores; SoRA (Ding et al., 2023) modulates the intrinsic rank during fine-tuning through proximal  
145 gradient-based gating; and HydraLoRA (Tian et al., 2024) introduces an asymmetric architecture that  
146 removes the need for task-specific tuning expertise in complex scenarios. Other works like (He et al.,  
147 2023a; Zhang et al., 2024b; Fu et al., 2022; Shen et al., 2024b), they focus on selecting or sparsifying  
148 subsets of existing parameters rather than learning sparse latent features.

149 Despite their efficiency, current PEFT approaches, whether using auxiliary modules as in prompt  
150 tuning and adapters, or low-rank approximations as in LoRA, share a common drawback: they  
151 encode adaptations in dense and entangled representations. This results in insufficient dynamics  
152 during the fine-tuning process, constrained representational capacity for multi-domain adaptation, and  
153 limited interpretability. As a result, there is growing interest in developing advanced PEFT methods  
154 that retain strong adaptation performance while introducing *structured sparsity*, with the goal of  
155 enhancing *fine-tuning dynamics*, *representational flexibility*, and *interpretability*.

156 **Sparse Autoencoder.** Sparse autoencoders (SAEs) (Ng et al., 2011; Makhzani & Frey, 2013)  
157 provide a scalable, unsupervised framework for learning compact and disentangled representations  
158 by imposing sparsity constraints on hidden activations. This methodology has recently gained  
159 substantial attention for its ability to enhance the interpretability of foundation models, including  
160 large language models (LLMs) (Cunningham et al., 2023; Templeton et al., 2024), vision-language  
161 models (VLMs) (Zhang et al., 2024a), and CLIP-style architectures (Lim et al., 2024). A core idea  
162 of SAEs lies in learning to reconstruct inputs using a sparse set of features in a higher-dimensional  
163 space, effectively disentangling superposed features (Briicken et al., 2023; Shi et al., 2025).

162 Specifically, given a data input  $x \in \mathbb{R}^n$ , encoder  $E \in \mathbb{R}^{l \times n}$ , and decoder  $D \in \mathbb{R}^{n \times l}$ , the autoencoding  
 163 process of a topk SAE (Gao et al., 2024) can be formulated as:

$$164 \quad z_k := \text{topk}(Ex), \quad \hat{x} := Dz_k,$$

166 where  $\text{topk}(\cdot)$  selects the top  $k$  ( $k \ll l$ ) largest elements of the input and zeroes out the remaining  
 167 values. This enforced sparsity constraint facilitates the emergence of semantically meaningful  
 168 representations  $z_k$ , effectively disentangling the complex, superposed features within large models.

169 Alongside interpretability, the potential of SAEs in other areas has also been explored. In efficiency  
 170 optimization, SAEs demonstrate remarkable versatility by addressing both retrieval and computational  
 171 constraints through sparse coding. Recent works (Kim et al., 2024; Wen et al., 2025) exemplify  
 172 this capability: the former optimizes the tradeoff between retrieval accuracy and efficiency through  
 173 sparse contrastive learning and sparse matrix factorization, while the latter compresses LLM key-  
 174 value caches using universal dictionaries. These approaches transform dense, high-dimensional  
 175 representations into sparse, efficient formats that preserve essential information while largely reducing  
 176 computational overheads.

177 In the context of safety alignment and controllable generation in diffusion models (Cywiński &  
 178 Deja, 2025; Kim & Ghadiyaram, 2025; Shi et al., 2025; Surkov et al., 2024; Tian et al., 2025),  
 179 SAEs are implemented as zero-shot concept detectors, enabling precise identification of both desired  
 180 and undesired features during the generative process. By isolating specific semantic features in the  
 181 model’s latent space, SAEs facilitate targeted interventions while maintaining the overall quality  
 182 of generation. Although SAEs have been adapted to a wide range of tasks and domains, their  
 183 benefits have been largely overlooked in the context of PEFT. This presents a significant research  
 184 gap: integrating the representational capabilities of SAEs with PEFT techniques can not only expand  
 185 the adaptation space, but also enable more flexible feature selection and improve interpretability.  
 186 In addition, the sparse structure of SAEs introduces greater dynamism into the adaptation process,  
 187 mitigating the rigidity typically associated with the low-rank spaces.

### 188 3 METHODOLOGY

191 **Background.** Parameter-efficient fine-tuning has become a key paradigm for adapting large-scale  
 192 pretrained foundation models to downstream tasks without incurring the full computational cost of  
 193 end-to-end fine-tuning (Xu et al., 2023). Among various PEFT methods, LoRA (Hu et al., 2022) has  
 194 emerged as one of the most widely used due to its simplicity and empirical effectiveness. Instead of  
 195 modifying the original weights directly, LoRA introduces a low-rank trainable update that captures  
 196 task-specific information.

197 Specifically, given a layer with a pretrained weight matrix  $W_0 \in \mathbb{R}^{m \times n}$ , LoRA models the task-  
 198 specific adaptation as a low-rank update to the weights:

$$199 \quad \Delta W = \frac{\alpha}{r} BA, \quad (1)$$

201 where  $A \in \mathbb{R}^{r \times n}$  and  $B \in \mathbb{R}^{m \times r}$  are low-rank matrices with  $r \ll \min(m, n)$ , and  $\alpha$  is a scaling  
 202 factor. The forward pass through the adapted layer, receiving input  $x \in \mathbb{R}^n$  and producing output  
 203  $h \in \mathbb{R}^m$ , is then modified as follows:

$$204 \quad h = W_0 x + \Delta W x. \quad (2)$$

206 Here,  $W_0$  remains frozen, and only the components representing  $\Delta W$  are updated during training.  
 207 The core idea of LoRA lies in constraining the update matrix  $\Delta W$  to possess a low intrinsic rank,  
 208 denoted by  $r$ . While this low-rank adaptation strategy is highly effective in terms of performance  
 209 and efficiency, compression of the task adaptation into dense, low-dimensional subspace defined by  
 210 matrices  $A$  and  $B$  inherently leads to the entanglement of potentially distinct underlying concepts,  
 211 posing significant challenges for interpretability and mechanistic understanding of learned adaptation.

212 However, while effective, LoRA has inherent limitations rooted in its architectural constraints. The  
 213 low-rank structure of the matrices  $A$  and  $B$  severely restricts the representation capacity of the  
 214 model (Wang & Zhao, 2025; Zhang et al., 2025), and the compression nature of LoRA reduces its  
 215 ability to capture *dynamic* patterns during parameter updates. This lack of adaptability may lead  
 to information obfuscation and hinders the model’s capacity to generalize across heterogeneous or

multidisciplinary domains. Furthermore, prior studies (Ding et al., 2023; He et al., 2025; Liu et al., 2023c; Mao et al., 2025; Zhang et al., 2023) have noted that compressing adaptations into a rigid low-rank subspace can result in a representational bottleneck, limiting the expressivity required for complex task-specific transformations. In addition, the dense structure of the low-rank projections often entangles distinct adaptation features, making it difficult to assign functional meaning to individual components within the update matrix  $\Delta W$ . These limitations underscore the necessity for a new PEFT paradigm that enables more adaptive and flexible representations while enabling dynamic, structured, and interpretable adaptation.

### 3.1 SPARSE ADAPTATION

To this end, we propose STAN, a high-dimensional, dynamic sparse adaptation framework, inspired by Sparse Autoencoders (Cunningham et al., 2023). It is designed to provide broader representational capacity and dynamic feature selection throughout the adaptation process. By introducing sparsity, STAN enhances the dynamics of the adaptation process, equipping the model with the capability to generalize across multiple domains and offering improved potential for interpretability.

Similar to LoRA, STAN modifies the forward pass by adding a learned adaptation  $\Delta Wx$  to the frozen pretrained output  $W_0x$ , as in Eq. 2 but changes how this adaptation is computed. Specifically, STAN introduces an encoder matrix  $E \in \mathbb{R}^{l \times n}$  and a decoder matrix  $D \in \mathbb{R}^{m \times l}$ , where  $l$  defines the dimension of a high-dimensional latent space. The input  $x$  is first projected into this space via  $E$ , and sparsity is enforced using a topk operator that selects only the top  $k$  activations, zeroing out the rest. The sparse code is then decoded by  $D$  to produce the adaptation signal:

$$\Delta Wx = \frac{1}{k} D \cdot \text{topk}(Ex). \quad (3)$$

The complete forward pass becomes:

$$h = W_0x + \frac{1}{k} D \cdot \text{topk}(Ex), \quad (4)$$

where  $k \ll l$  is a sparsity-controlling hyperparameter. Only  $E$  and  $D$  are updated during training, while  $W_0$  remains frozen.

### 3.2 STAN AS A NON-LINEAR COMBINATION OF SUBSPACES

A **key distinction** between STAN and LoRA lies in the fact that STAN’s adaptation  $\Delta Wx$  is a nonlinear, input-dependent function. This nonlinearity stems from the topk operator, which *dynamically* selects a different subset of latent dimensions depending on  $x$ .

Let  $\mathcal{I}(x) \subset 1, \dots, l$  denote the indices of the top- $k$  activations for a given input  $x$ . We define an input-dependent selection matrix  $S(x) \in \{0, 1\}^{l \times l}$  as a diagonal matrix with  $S(x)_{ii} = 1$  if  $i \in \mathcal{I}(x)$  and 0 otherwise. This allows the sparsity operation to be expressed as  $\text{topk}(Ax) = S(x)Ax$ . Substituting into Equation 3 yields:

$$\Delta Wx = DS(x)Ex. \quad (5)$$

Equivalently, let  $E_{\mathcal{I}(x)} \in \mathbb{R}^{k \times n}$  denote the submatrix of  $E$  consisting of the rows indexed by  $\mathcal{I}(x)$ , and let  $D_{\mathcal{I}(x)} \in \mathbb{R}^{m \times k}$  denote the submatrix of  $D$  consisting of the corresponding columns. Then, for a given input  $x$ , the adaptation is equivalent to a projection onto a  $k$ -dimensional subspace defined by the selected latent features:

$$\Delta Wx = D_{\mathcal{I}(x)}E_{\mathcal{I}(x)}x. \quad (6)$$

Since the index set  $\mathcal{I}(x)$  varies with the input  $x$ , STAN effectively operates over a mixture of input-dependent  $k$ -dimensional subspaces, each spanned by a distinct subset of the  $l$ -dimensional latent feature space. The topk operator functions as a dynamic, non-linear router, selecting the most relevant features for each input. This mechanism enables STAN to model a richer class of adaptation functions than LoRA, which is confined to a single, fixed low-rank subspace. By *dynamically* composing basis vectors from a larger representational space, STAN offers greater expressivity while maintaining sparsity. Furthermore, in contrast to LoRA, where increasing the rank results in an exponential growth in trainable parameters, the large representational capacity offered by STAN allows for *flexible scaling* of feature selection without increasing the overall number of training parameters. For more discussion with related methods, see Appendix B.

In summary, the **core** of STAN is to leverage a non-linear, input-dependent sparse activation mechanism within a high-dimensional latent space to *dynamically* compose multiple adaptation subspaces, aiming for both representational flexibility and interpretability through sparsity. For more comprehensive discussion towards the interpretability and identifiability of STAN, please refer to Appendix A.

Model	Method	Accuracy ↑				Matthew's Corr. ↑
		MNLI	SST-2	QNLI	QQP	
RoBERTa-base	LoRA	0.8514	0.9177	0.9177	0.8627	0.5981
	AdaLoRA	0.8429	0.9358	0.9225	0.8812	0.6132
	SoRA	0.7657	0.9220	0.8380	0.8420	0.5485
	STAN (Ours)	<b>0.9303</b>	<b>0.9495</b>	<b>0.9408</b>	<b>0.9242</b>	<b>0.6191</b>
RoBERTa-large	LoRA	0.8812	0.9553	0.9131	0.8842	0.6749
	AdaLoRA	0.8857	0.9472	0.9400	0.8883	0.6314
	SoRA	0.8769	0.9280	0.4860	0.8450	0.3470
	STAN (Ours)	<b>0.8919</b>	<b>0.9610</b>	<b>0.9489</b>	<b>0.8957</b>	<b>0.7400</b>
DeBERTaV3-base	LoRA	0.8857	0.9438	0.9371	0.9163	0.6729
	AdaLoRA	0.8637	0.9553	0.9440	0.8952	0.6864
	SoRA	0.8095	0.9564	0.9322	0.8540	0.6698
	STAN (Ours)	<b>0.8974</b>	<b>0.9622</b>	<b>0.9477</b>	<b>0.9230</b>	<b>0.6904</b>
DeBERTaV3-large	LoRA	0.8879	0.9599	0.9503	0.8923	0.7237
	AdaLoRA	0.9021	0.9587	0.9552	0.8899	0.7008
	SoRA	0.9056	0.9370	0.9440	0.8640	0.6829
	STAN (Ours)	<b>0.9145</b>	<b>0.9622</b>	<b>0.9590</b>	<b>0.9058</b>	<b>0.7528</b>

Table 1: Performance comparison on language understanding tasks across four large pretrained language models using five tasks from the GLUE benchmark. ↑ indicates that higher values are better. The best results are highlighted in **bold**.

Method	QNLI	MNLI	SST-2	QQP	MRPC	RTE	STSB
LoRA	0.9371	0.8857	0.9438	0.9163	0.8995	0.8520	0.9160
AdaLoRA	0.9440	0.8637	0.9553	0.8952	0.9069	0.8736	0.9163
SoRA	0.9322	0.8095	0.9564	0.8540	0.8734	0.8777	<u>0.9222</u>
PiSSA	<u>0.9443</u>	0.8729	0.9621	<u>0.9230</u>	<u>0.9150</u>	0.8869	0.9200
BOFT	0.9423	<b>0.9025</b>	<b>0.9644</b>	0.9210	0.9016	<u>0.8881</u>	0.9192
SVFT(P)	0.9427	0.8969	0.9541	0.9016	0.8877	0.8724	0.9180
SVFT(R)	0.9390	0.8805	0.9002	0.9150	0.8899	0.8809	0.9173
VeRA	0.9324	<u>0.8993</u>	0.9553	0.9040	0.8794	0.8700	0.8871
LoRA-XS	0.8475	0.7802	0.9243	0.8038	0.8132	0.8065	0.8231
LS-LoRA	0.9235	0.8736	0.9427	0.8757	0.8382	0.6751	0.8767
LoRETTA	0.9325	0.8680	0.9553	0.8920	0.8873	0.7581	0.9066
STAN (Ours)	<b>0.9477</b>	0.8974	<u>0.9622</u>	<b>0.9230</b>	<b>0.9166</b>	<b>0.9114</b>	<b>0.9277</b>

Table 2: Results with DeBERTaV3-base across more baseline methods. The best results are highlighted in **bold**, and the second best results are underlined. All are measured by accuracy.

## 4 EXPERIMENTS

In this section, we conduct a series of experiments to validate the effectiveness of our proposed method, STAN. In Section 4.1, we evaluate its task performance across a range of benchmarks and model architectures, comparing it against state-of-the-art and representative PEFT methods to demonstrate its efficacy. In Section 4.2, we further explore the applicability of STAN to the Stable Diffusion 3 (SD3) model by fine-tuning it on a style alignment task, assessing its impact on both generation quality and adaptation flexibility. We also present qualitative visualizations that highlight the disentangled features learned by STAN, offering insight into its representational behavior. And we place more experiments related to interpretability in Appendix H. Finally, Section 4.3 presents ablation studies examining the role of sparsity levels and their influence on both performance and representation quality and in Section 4.4 we studies about the catastrophic forgetting phenomena comparing to LoRA. Appendix D analyzes the training dynamics and convergence behavior of STAN, focusing on its stability and efficiency during fine-tuning. All results are reproduced by us in a unified training environment using official implementations.

324	325	326	327	Dataset	LLaVA-1.5-7B		LLaVA-1.5-13B		LLaVA-1.6-Vicuna-13B	
					LoRA	STAN (Ours)	LoRA	STAN (Ours)	LoRA	STAN (Ours)
				GQA	80.60	<b>82.36</b>	81.49	<b>83.16</b>	82.29	<b>83.77</b>
				ScienceQA	89.17	<b>90.25</b>	91.34	<b>92.96</b>	92.06	<b>92.49</b>

328 Table 3: Performance comparison on multimodal (vision-language) tasks across three LLaVA model  
329 variants using the GQA and ScienceQA benchmarks. The best results are highlighted in **bold**.

331	Method	GSM8K	MATH	Avg.
332	LoRA	60.6	10.8	35.7
333	PiSSA	58.2	10.4	34.3
334	STAN (Ours)	<b>60.7</b>	<b>11.6</b>	<b>36.15</b>

335 Table 4: Performance on math & code benchmarks on LLaMA-2-7B. Best is in **bold**.

331	Method	VizWiz	POPE	MMBench	Avg.
332	LoRA	0.5021	0.8549	0.5403	0.6324
333	STAN (Ours)	<b>0.5109</b>	<b>0.8611</b>	<b>0.5416</b>	<b>0.6379</b>

334 Table 5: Performance on multimodal benchmarks on LLaVA-  
335 1.5-7B. The best results are highlighted in **bold**.

#### 337 4.1 QUANTITATIVE COMPARISON

339 In this section, we compare our method against baseline approaches on single-modality (language),  
340 reasoning (math & code) and multi-modality (vision-language) benchmarks to demonstrate the  
341 validity and advantages of STAN. For language tasks, we compare STAN with LoRA (Hu et al.,  
342 2022), AdaLoRA (Zhang et al., 2023), and SoRA (Ding et al., 2023), using four pretrained large  
343 language models: RoBERTa-base/large (Liu et al., 2019) and DeBERTaV3-base/large (He et al.,  
344 2023b). These models are evaluated on five tasks from the GLUE benchmark (Wang et al., 2018).  
345 We report accuracy for MNLI, SST-2, QNLI, and QQP, and use Matthew’s correlation for CoLA.  
346 Further, we present a more comprehensive comparison experiment with wider range of baseline  
347 methods (Hu et al., 2022; Zhang et al., 2023; Ding et al., 2023; Meng et al., 2024; Liu et al., 2023d;  
348 Lingam et al., 2024; Kopitzko et al., 2023; Bałazy et al., 2024; He et al., 2022; Yang et al., 2024) on  
349 DeBERTaV3-base (He et al., 2023b) with GLUE benchmark (Wang et al., 2018). The results are  
350 summarized in Table 1 and Table 2. For multimodal tasks, we evaluate STAN against LoRA on three  
351 variants of the LLaVA model (Liu et al., 2023a), using the GQA (Hudson & Manning, 2019) and  
352 ScienceQA (Lu et al., 2022) benchmarks. Moreover, we compare STAN against LoRA using more  
353 complex benchmarks, VizWiz (Gurari et al., 2018), POPE (Li et al., 2023) and MMBench (Liu et al.,  
354 2024b), on LLaVA-1.5-7b (Liu et al., 2023a), to further demonstrate the robustness of our method. All  
355 benchmarks are evaluated using accuracy and the results are summarized in Table 3 and Table 5. As  
356 for the reasoning task, we compare our STAN method with LoRA (Hu et al., 2022) and PiSSA (Meng  
357 et al., 2024) on GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) benchmarks using  
358 LLaMA-2-7B (Touvron et al., 2023). The results are shown on Table 4. Experimental settings and  
359 runtime analysis are detailed in Appendix C.

360 As shown in the above Tables, the experimental results consistently demonstrate the superior performance  
361 and stability of our proposed method across a diverse range of tasks and model architectures.  
362 On language understanding benchmarks, STAN frequently achieves the highest scores across all  
363 five datasets when applied to different models of varying sizes and structures, often outperforming  
364 established PEFT methods by a noticeable margin. Notably, although SoRA (Ding et al., 2023) also  
365 incorporates sparsity, it exhibits significant instability during fine-tuning, whereas STAN delivers  
366 consistently strong performance (See more details in Appendix D). This pattern of performance  
367 extends to reasoning and vision-language tasks, where STAN consistently outperforms other methods  
368 across all benchmarks with different models. The consistent top-tier results across modalities, model  
369 types, model scales, and benchmarks highlight not only the enhanced performance of STAN, but also  
370 its robust stability and generalizability.

#### 371 4.2 STYLE ALIGNMENT

372 In this section, we evaluate the capability of STAN to perform style alignment in the context of image  
373 generation. Our experiments are conducted on Stable Diffusion 3 (SD3) (Rombach et al., 2022), with  
374 a mixed dataset comprising images from WikiArt (Saleh & Elgammal, 2015) and the DualStyleGAN  
375 dataset (Yang et al., 2022), offering a diverse range of artistic styles.

376 We compare the performance of three approaches: the pretrained SD3 model without any fine-tuning  
377 (denoted as None), the SD3 model fine-tuned using LoRA, and the SD3 model fine-tuned using  
378 our proposed STAN method. To quantitatively evaluate style alignment, we employ two metrics:



Figure 2: Qualitative demonstration of style alignment results across different methods using SD3. None denotes the pretrained model without fine-tuning.

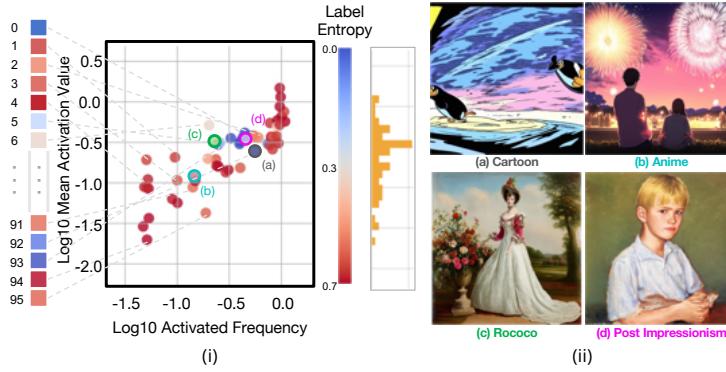


Figure 3: Qualitative demonstration of the sparse representations learned by STAN. (i) Distribution of sparsely activated intermediate features across four distinct styles. (ii) Visualizations corresponding to each style: (a) Cartoon, (b) Anime, (c) Rococo, and (d) Post Impressionism.

CLIP-Score (Radford et al., 2021) and DINO-Score (Caron et al., 2021), both of which measure the semantic similarity between the generated image and the corresponding textual style prompt, but rely on different pretrained vision-language models. As shown in Table 6, STAN achieves the highest performance on both metrics, with a CLIP-Score of 0.6694 and a DINO-Score of 0.4283, outperforming both the baseline and the LoRA-fine-tuned variant.

To ensure that the observed quantitative improvements are also perceptually meaningful, we conducted an additional human user study focusing on both style alignment fidelity and overall visual quality in the SD3 experiments. We randomly sampled 20 distinct style prompts, and for each prompt, participants were presented with three images generated by the None, LoRA, and STAN methods with randomized orders. Users were asked to answer the question: Which image best matches the target style and has the highest overall visual quality?

The aggregated win rates, detailed in Table 7, strongly corroborate the preceding CLIP and DINO scores. STAN is preferred in over 90% of the comparisons, demonstrating a significant superiority in both style alignment and perceived image quality over the pretrained baseline and the LoRA-finetuned model. These results validate that the performance gains achieved by STAN are not merely numerical artifacts but translate to substantial enhancements in the human perception of the generated content. Combined with the disentanglement analysis in our style experiments, these findings further strengthen the claim that STAN learns disentangled, style-specific sparse features, supporting more precise and controllable generation.

Methods	CLIP-Score↑
None	0.5572
LoRA	0.6509
STAN (Ours)	<b>0.6694</b>
Methods	DINO-Score↑
None	0.3383
LoRA	0.4142
STAN (Ours)	<b>0.4283</b>

Table 6: Quantitative comparison of style alignment results using SD3. Best in **bold**.

432 The superior quantitative results of STAN are further supported by vi-  
 433 sual evidence, as illustrated in Figure 2. The sparse latent space induced  
 434 by STAN facilitates more effective disentanglement and encoding of  
 435 stylistic elements, enabling more accurate alignment with the intended  
 436 artistic styles. Visualizations show that the pretrained model often fails  
 437 to capture or render the specified styles, while LoRA improves align-  
 438 ment but may still result in stylistic confusion. For example, blending  
 439 characteristics of Realism and Impressionism, or generate inconsistent  
 440 outputs in styles such as Cartoon or Illustration. In contrast, images produced by the STAN-fine-tuned model exhibit the most  
 441 faithful and distinct adherence to the target styles, demonstrating its improved ability to isolate and  
 442 apply diverse artistic features. These results highlight not only the enhanced representational capacity  
 443 of STAN, but also its ability to decouple semantically similar yet stylistically distinct information.  
 444 Additional visual examples are provided in Appendix E.

445 Furthermore, to highlight disentangled features learned by STAN, we present a statistical analysis in  
 446 Figure 3. Part (i) illustrates activation frequency statistics of sparsely activated intermediate features  
 447 when the model is prompted with four distinct artistic styles, i.e., Cartoon, Anime, Rococo, and Post  
 448 Impressionism, using 300 generated samples per style. The y-axis represents  $\log_{10}$  mean activation  
 449 value, and the x-axis indicates  $\log_{10}$  activation frequency. Each point is colored according to its label  
 450 entropy, which reflects the degree of style specificity associated with that neuron: lower entropy  
 451 values indicate specialization (i.e., the neuron is primarily activated by one style), while higher  
 452 entropy suggests activation by a broader mix of styles. Details on the formulation of this analysis  
 453 can be found in Appendix F. Part (ii) presents corresponding visual examples, illustrating that the  
 454 images activating a given neuron are consistently aligned with its associated artistic style. The results  
 455 reveal that different styles tend to activate distinct, often non-overlapping, subsets of sparse latent  
 456 features. A small number of shared neurons appear to capture common generative priors, while the  
 457 majority remain style-specific. This separation in feature activations suggests that STAN is effectively  
 458 associating specific sparse components with semantically coherent stylistic concepts. Moreover, it  
 459 demonstrates the expressive capacity of subspace combinations in the sparse latent space. To better  
 460 demonstrate the interpretability of STAN, we did the similar experiments on the language task as  
 461 well. Please refer to Appendix H for more details.

462 We also conduct additional experiments where we apply LoRA under the  
 463 same setup and probe its latent directions using the same attribution and  
 464 activation-frequency analysis. As demonstrated in Figure 4, The dense low-  
 465 rank adaptations lead to more entangled features compared to STAN’s sparse  
 466 latent units, and we can not separate styles according to the middle features  
 467 like we did to STAN.

468 The ability of STAN to map these diverse styles to separable sparse feature  
 469 activations is a strong indication of its capacity for learning disentangled  
 470 features. Compared to LoRA, this broader representational capacity is essential  
 471 for achieving fine-grained control in generative modeling.

### 472 4.3 ABLATION STUDY

473 In this section, we conduct the ablations to investigate the impact of the  $TopK$  parameter along with  
 474 the middle latent dimension  $l$  on the performance of our STAN method. The evaluation is based  
 475 on accuracy on a subset of the MNLI dataset (Wang et al. 2018). We control the size of the latent  
 476 space using an expansion factor, which scales the dimensionality of the intermediate representation  
 477 by multiplying it with the input feature dimension. The  $TopK$  parameter controls the sparsity level by  
 478 selecting the  $K$  most active features in this latent space. In our experiments, we vary the expansion  
 479 factor across 1, 1.5, 2 and  $TopK$  across 1, 4, 8, 16, 32, 64, along with a setting denoted as  $n$ -latents,  
 480 where all intermediate features are retained (i.e., no sparsification is applied). The full results of the  
 481 ablation study are presented in Table 8.

482 As in Table 8, we can see a nuanced interplay between these two hyperparameters. For all three  
 483 expansion factors, accuracy peaks at  $K=8$  or  $K=16$ , which indicates that as the latent space capacity  
 484 increases, a slightly larger number of active features might be beneficial, but it does not mean the  
 485 elevation of active feature size will bring the better performance.

Method	Wins	Total	Win Rate (%)
None	6	501	1.20
LoRA	39	501	7.78
Ours	456	501	91.02

Table 7: Human Evaluation of Style Alignment and Image Quality.

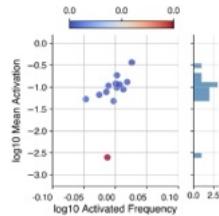


Figure 4: Interpretability experiment using LoRA.

486 Another finding emerges when comparing the sparse Top $K$  config-  
 487 urations with the  $n$ -latents baseline. Even when the it configuration  
 488 utilizes a substantially larger, high-rank intermediate dimension,  
 489 with expansion factors equal to 1.5 and 2, achieving accuracies  
 490 of 0.8262 and 0.8286 respectively, STAN configurations with a  
 491 small Top $K$  value can achieve remarkably comparable results. This  
 492 implies that a significant portion of the features within the larger,  
 493 unparsified latent space may be redundant for the given task. The  
 494 superior performance of STAN with small  $K$  values, even when  
 495 the potential latent dimension is large, underscores the efficacy of  
 496 sparsity, it allows the model to achieve competitive results by iden-  
 497 tifying a compact set of the most salient features, thereby alleviating  
 498 the need for an excessively large number of active parameters and demonstrating that sparsity is a  
 499 sensible and efficient approach to parameter utilization.

#### 4.4 CATASTROPHIC FORGETTING

500  
 501 To directly evaluate the catastrophic forgetting, we conducted additional bidirectional sequential  
 502 finetuning experiments on two tasks: SST-2 and MNLI. We perform sequential finetuning in two  
 503 directions: SST-2  $\rightarrow$  MNLI: Finetune using STAN on SST-2 first, then continue training on MNLI  
 504 while monitoring SST-2 accuracy. MNLI  $\rightarrow$  SST-2: Finetune on MNLI, then finetune on SST-2 while  
 505 monitoring MNLI accuracy. This setup directly measures forgetting of the first task during learning  
 506 of the second. Table 9 demonstrates the experiment results.

507  
 508 We can see that for SST-2  $\rightarrow$  MNLI, only a 4.6-point drop  
 509 after 20 epochs, and for MNLI  $\rightarrow$  SST-2, it's a moderate  
 510 6.5-point decline. The results indicate a very mild forget-  
 511 ting. We also performed LoRA based experiments utilizing  
 512 the same bidirectional setting, its drop for SST-2  $\rightarrow$  MNLI  
 513 is 14.2 points after 20 epochs and for MNLI  $\rightarrow$  SST-2 the  
 514 drop is 17.8 points. That is because STAN's architecture  
 515 naturally mitigates forgetting due to (1) sparse latent ac-  
 516 tivations that localize task-specific updates, (2) minimal  
 517 interference across tasks, since only a small subset of neu-  
 518 rons is modified. (3) disentangled and interpretable features, allowing different tasks to occupy  
 519 different sparse subspaces. This is fundamentally different from LoRA's dense low-rank updates,  
 520 which modify shared directions and are more prone to overwriting previous knowledge.

## 5 CONCLUSION AND FUTURE WORKS

521 In this paper, we introduced STAN, a novel parameter-efficient fine-tuning method designed to  
 522 support dynamic feature selection and address key limitations of existing methods – particularly the  
 523 restricted representational capacity inherent in low-rank adaptation paradigms such as LoRA. By  
 524 leveraging a sparse autoencoder architecture, STAN dynamically encodes task-specific adaptations  
 525 through a high-dimensional yet sparse set of features. This enables a more flexible and expressive  
 526 adaptation mechanism compared to rigid low-rank projections. For future work, STAN can be  
 527 extended to a broader range of tasks and modalities, including more complex reasoning and multi-  
 528 step decision-making settings. In addition, further direction can be explored, such as investigating  
 529 alternative sparsity-inducing mechanisms beyond the current Top $K$  selection strategy. For the use of  
 530 large language models, please refer to Appendix I.

531 **Ethics statement.** We have adhered to the ICLR Code of Ethics. Our research primarily utilizes  
 532 publicly available datasets and pretrained models, and we do not foresee any direct negative societal  
 533 impacts or ethical concerns arising from our methodology.

534 **Reproducibility statement.** We are committed to ensuring the full reproducibility of our research.  
 535 To facilitate this, the complete source code to replicate experiments presented in this paper will be  
 536 made publicly available upon publication.

Top $K$	Expansion Factor		
	1	1.5	2
1	0.8139	0.8088	0.8131
4	0.8158	0.8122	0.8149
8	<b>0.8192</b>	<u>0.8206</u>	0.8177
16	0.8171	0.8180	<u>0.8213</u>
32	<u>0.8173</u>	0.8170	0.8169
64	0.8162	0.8155	0.8161
$n$ -latents	0.8158	<b>0.8262</b>	<u>0.8286</u>

Table 8: Ablation study on expansion factor and Top $K$ . Best is in **bold**, and second is in underlined.

method	Accuracy (%) at Epoch				
	1	5	10	15	20
SST-2	85.00	84.03	82.83	81.62	80.41
MNLI	80.00	78.62	76.90	75.17	73.45

Table 9: Catastrophic forgetting study using STAN. The first row is the accuracy change of SST-2 in SST-2  $\rightarrow$  MNLI setting the the second row is the accuracy change of MNLI in MNLI  $\rightarrow$  SST-2.

540 REFERENCES  
541

542 Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge,  
543 Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.

544 Klaudia Bałazy, Mohammadreza Banaei, Karl Aberer, and Jacek Tabor. Lora-xs: Low-rank adaptation  
545 with extremely small number of parameters. *arXiv preprint arXiv:2405.17604*, 2024.

546

547 Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick  
548 Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec,  
549 Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina  
550 Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and  
551 Christopher Olah. Towards monosemanticity: Decomposing language models with dictionary  
552 learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.

553

554 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
555 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
556 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

557

558 Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer:  
559 Evaluating and testing unintended memorization in neural networks. In *28th USENIX security  
560 symposium (USENIX security 19)*, pp. 267–284, 2019.

561

562 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and  
563 Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the  
564 International Conference on Computer Vision (ICCV)*, 2021.

565

566 Guanhua Chen, Yutong Yao, Ci-Jun Gao, Lidia S Chao, Feng Wan, and Derek F Wong. Not all lora  
567 parameters are essential: Insights on inference necessity. *arXiv preprint arXiv:2503.23360*, 2025.

568

569 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
570 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve  
math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

571

572 Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoen-  
573 coders find highly interpretable features in language models. *arXiv preprint arXiv:2309.08600*,  
2023.

574

575 Bartosz Cywiński and Kamil Deja. Saeuron: Interpretable concept unlearning in diffusion models  
576 with sparse autoencoders. *arXiv preprint arXiv:2501.18052*, 2025.

577

578 Ning Ding, Xingtai Lv, Qiaosen Wang, Yulin Chen, Bowen Zhou, Zhiyuan Liu, and Maosong Sun.  
579 Sparse low-rank adaptation of pre-trained language models. *arXiv preprint arXiv:2311.11696*,  
2023.

580

581 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
582 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An  
583 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint  
584 arXiv:2010.11929*, 2020.

585

586 Zihao Fu, Haoran Yang, Anthony Man-Cho So, Wai Lam, Lidong Bing, and Nigel Collier. On  
587 the effectiveness of parameter-efficient fine-tuning, 2022. URL <https://arxiv.org/abs/2211.15583>.

588

589 Leo Gao, Tom Dupré la Tour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya  
590 Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. *arXiv preprint  
591 arXiv:2406.04093*, 2024.

592

593 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

594 Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and  
 595 Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In  
 596 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3608–3617,  
 597 2018.

598 Haoyu He, Jianfei Cai, Jing Zhang, Dacheng Tao, and Bohan Zhuang. Sensitivity-aware visual  
 599 parameter-efficient fine-tuning. In *ICCV*, 2023a.

601 Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a  
 602 unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*, 2021.

603

604 Pengcheng He, Jianfeng Gao, and Weizhu Chen. DeBERTaV3: Improving deBERTa using ELECTRA-  
 605 style pre-training with gradient-disentangled embedding sharing. In *The Eleventh International  
 606 Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=sE7-XhLxHA>.

607

608 Shuai He, Liang Ding, Daize Dong, Miao Zhang, and Dacheng Tao. Sparseadapter: An easy  
 609 approach for improving the parameter-efficiency of adapters. *arXiv preprint arXiv:2210.04284*,  
 610 2022.

611

612 Zhiwei He, Zhaopeng Tu, Xing Wang, Xingyu Chen, Zhijie Wang, Jiahao Xu, Tian Liang, Wenxiang  
 613 Jiao, Zhuosheng Zhang, and Rui Wang. RaSA: Rank-sharing low-rank adaptation. In *The Thirteenth  
 614 International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=GdXI5zCoAt>.

615

616 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
 617 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv  
 618 preprint arXiv:2103.03874*, 2021.

619

620 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in  
 621 neural information processing systems*, 33:6840–6851, 2020.

622

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

626

627 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
 628 and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International  
 629 Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.

630

631 Zhiqiang Hu, Lei Wang, Yihuai Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya  
 632 Poria, and Roy Ka-Wei Lee. Llm-adapters: An adapter family for parameter-efficient fine-tuning  
 633 of large language models. *arXiv preprint arXiv:2304.01933*, 2023.

634

635 Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning  
 636 and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer  
 vision and pattern recognition*, pp. 6700–6709, 2019.

637

638 Dahye Kim and Deepti Ghadiyaram. Concept steerers: Leveraging k-sparse autoencoders for  
 639 controllable generations. *arXiv preprint arXiv:2501.19066*, 2025.

640

641 Junhyuck Kim, Jongho Park, Jaewoong Cho, and Dimitris Papailiopoulos. Lexico: Extreme kv cache  
 642 compression via sparse coding over universal dictionaries, 2024. URL <https://arxiv.org/abs/2412.08890>.

643

644 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
 645 language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:  
 646 22199–22213, 2022.

647

648 Dawid J Kopiczko, Tijmen Blankevoort, and Yuki M Asano. Vera: Vector-based random matrix  
 649 adaptation. *arXiv preprint arXiv:2310.11454*, 2023.

648 Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt  
 649 tuning. *arXiv preprint arXiv:2104.08691*, 2021.  
 650

651 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv*  
 652 *preprint arXiv:2101.00190*, 2021.

653 Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object  
 654 hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023.  
 655

656 Hyesu Lim, Jinho Choi, Jaegul Choo, and Steffen Schneider. Sparse autoencoders reveal selective  
 657 remapping of visual concepts during adaptation. *arXiv preprint arXiv:2412.05276*, 2024.

658 Vijay Chandra Lingam, Atula Neerkaje, Aditya Vavre, Aneesh Shetty, Gautham Krishna Gudur,  
 659 Joydeep Ghosh, Eunsol Choi, Alex Dimakis, Aleksandar Bojchevski, and Sujay Sanghavi. Svft:  
 660 Parameter-efficient fine-tuning with singular vectors. *Advances in Neural Information Processing*  
 661 *Systems*, 37:41425–41446, 2024.

662

663 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023a.

664

665 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in*  
 666 *neural information processing systems*, 36:34892–34916, 2023b.

667

668 Qidong Liu, Xian Wu, Xiangyu Zhao, Yuanshao Zhu, Derong Xu, Feng Tian, and Yefeng Zheng.  
 669 Moelora: An moe-based parameter efficient fine-tuning method for multi-task medical applications.  
*CoRR*, 2023c.

670

671 Weiyang Liu, Zeju Qiu, Yao Feng, Yuliang Xiu, Yuxuan Xue, Longhui Yu, Haiwen Feng, Zhen  
 672 Liu, Juyeon Heo, Songyou Peng, et al. Parameter-efficient orthogonal finetuning via butterfly  
 673 factorization. *arXiv preprint arXiv:2311.06243*, 2023d.

674

675 Xinyang Liu, Dongsheng Wang, Bowei Fang, Miaoge Li, Yishi Xu, Zhibin Duan, Bo Chen, and  
 676 Mingyuan Zhou. Patch-prompt aligned bayesian prompt tuning for vision-language models. In  
*Uncertainty in Artificial Intelligence*, pp. 2309–2330. PMLR, 2024a.

677

678 Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike  
 679 Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining  
 680 approach. *arXiv preprint arXiv:1907.11692*, 2019.

681

682 Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi  
 683 Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player?  
 In *European conference on computer vision*, pp. 216–233. Springer, 2024b.

684

685 Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord,  
 686 Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for  
 687 science question answering. In *The 36th Conference on Neural Information Processing Systems*  
*(NeurIPS)*, 2022.

688

689 Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. An empirical study of  
 690 catastrophic forgetting in large language models during continual fine-tuning. *arXiv preprint*  
*arXiv:2308.08747*, 2023.

691

692 Alireza Makhzani and Brendan Frey. K-sparse autoencoders. *arXiv preprint arXiv:1312.5663*, 2013.

693

694 Yuren Mao, Yuhang Ge, Yijiang Fan, Wenyi Xu, Yu Mi, Zhonghao Hu, and Yunjun Gao. A survey  
 695 on lora of large language models. *Frontiers of Computer Science*, 19(7):197605, 2025.

696

697 Fanxu Meng, Zhaojun Wang, and Muhan Zhang. Pissa: Principal singular values and singular vectors  
 698 adaptation of large language models. *Advances in Neural Information Processing Systems*, 37:  
 121038–121072, 2024.

699

700 Fuseini Mumuni and Alhassan Mumuni. Explainable artificial intelligence (xai): from inherent  
 701 explainability to large language models. *arXiv preprint arXiv:2501.09967*, 2025.

Andrew Ng et al. Sparse autoencoder. *CS294A Lecture notes*, 72(2011):1–19, 2011.

702 Atharva Nijasure, Tanya Chowdhury, and James Allan. How relevance emerges: Interpreting lora  
 703 fine-tuning in reranking llms. *arXiv preprint arXiv:2504.08780*, 2025.  
 704

705 Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy  
 706 employed by v1? *Vision research*, 37(23):3311–3325, 1997.  
 707

708 Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder,  
 709 Kyunghyun Cho, and Iryna Gurevych. Adapterhub: A framework for adapting transformers. *arXiv*  
 710 *preprint arXiv:2007.07779*, 2020.  
 711

712 Nusrat Jahan Prrottasha, Upama Roy Chowdhury, Shetu Mohanto, Tasfia Nuzhat, Abdullah As Sami,  
 713 Md Shamol Ali, Md Shohanur Islam Sobuj, Hafijur Raman, Md Kowsher, and Ozlem Ozmen  
 714 Garibay. Peft a2z: Parameter-efficient fine-tuning survey for large language and vision models.  
 715 *arXiv preprint arXiv:2504.14117*, 2025.  
 716

717 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 718 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
 719 models from natural language supervision. In *International conference on machine learning*, pp.  
 8748–8763. PMLR, 2021.  
 720

721 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 722 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-  
 723 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.  
 724

725 Babak Saleh and Ahmed Elgammal. Large-scale classification of fine-art paintings: Learning the  
 726 right metric on the right feature. *arXiv preprint arXiv:1505.00855*, 2015.  
 727

728 Sheng Shen, Shijia Yang, Tianjun Zhang, Bohan Zhai, Joseph E Gonzalez, Kurt Keutzer, and Trevor  
 729 Darrell. Multitask vision-language prompt tuning. In *Proceedings of the IEEE/CVF Winter  
 730 Conference on Applications of Computer Vision*, pp. 5656–5667, 2024a.  
 731

732 Shufan Shen, Junshu Sun, Xiangyang Ji, Qingming Huang, and Shuhui Wang. Ex-  
 733 panding sparse tuning for low memory usage. In A. Globerson, L. Mackey, D. Bel-  
 734 grave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural In-  
 735 formation Processing Systems*, volume 37, pp. 76616–76642. Curran Associates, Inc.,  
 736 2024b. URL [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/8c420176b45e923cf99deel7356a763-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/8c420176b45e923cf99deel7356a763-Paper-Conference.pdf).  
 737

738 Yingdong Shi, Changming Li, Yifan Wang, Yongxiang Zhao, Anqi Pang, Sibei Yang, Jingyi Yu, and  
 739 Kan Ren. Dissecting and mitigating diffusion bias via mechanistic interpretability. In *Proceedings  
 740 of the Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 8192–8202, June 2025.  
 741

742 Viacheslav Surkov, Chris Wendler, Mikhail Terekhov, Justin Deschenaux, Robert West, and Caglar  
 743 Gulcehre. Unpacking sdxl turbo: Interpreting text-to-image models with sparse autoencoders.  
 744 *arXiv preprint arXiv:2410.22366*, 2024.  
 745

746 Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen,  
 747 Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L  
 748 Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers,  
 749 Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan.  
 750 Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Trans-  
 751 former Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.  
 752

753 Chunlin Tian, Zhan Shi, Zhijiang Guo, Li Li, and Cheng-Zhong Xu. Hydralora: An asymmetric  
 754 lora architecture for efficient fine-tuning. *Advances in Neural Information Processing Systems*, 37:  
 755 9565–9584, 2024.  
 756

757 Zhihua Tian, Sirun Nan, Ming Xu, Shengfang Zhai, Wenjie Qu, Jian Liu, Kui Ren, Ruoxi Jia, and  
 758 Jiaheng Zhang. Sparse autoencoder as a zero-shot classifier for concept erasing in text-to-image  
 759 diffusion models, 2025. URL <https://arxiv.org/abs/2503.09446>.  
 760

756 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay  
 757 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation  
 758 and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

759

760 Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE:  
 761 A multi-task benchmark and analysis platform for natural language understanding. In Tal Linzen,  
 762 Grzegorz Chrupała, and Afra Alishahi (eds.), *Proceedings of the 2018 EMNLP Workshop Black-*  
 763 *boxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, Brussels, Belgium,  
 764 November 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5446. URL  
 765 <https://aclanthology.org/W18-5446/>.

766

767 Maolin Wang and Xiangyu Zhao. Metalora: Tensor-enhanced adaptive low-rank fine-tuning. *arXiv*  
 768 preprint *arXiv:2504.00460*, 2025.

769

770 Yifan Wang, Yifei Liu, Yingdong Shi, Changming Li, Anqi Pang, Sibei Yang, Jingyi Yu, and Kan  
 771 Ren. Discovering influential neuron path in vision transformers. In *The Thirteenth International*  
 772 *Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=WQQyJbr5Lh>.

773

774 Tiansheng Wen, Yifei Wang, Zequn Zeng, Zhong Peng, Yudi Su, Xinyang Liu, Bo Chen, Hongwei  
 775 Liu, Stefanie Jegelka, and Chenyu You. Beyond matryoshka: Revisiting sparse coding for adaptive  
 776 representation, 2025. URL <https://arxiv.org/abs/2503.01776>.

777

778 Wenhan Xia, Chengwei Qin, and Elad Hazan. Chain of lora: Efficient fine-tuning of language models  
 779 via residual learning. *arXiv preprint arXiv:2401.04151*, 2024.

780

781 Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. Parameter-efficient  
 782 fine-tuning methods for pretrained language models: A critical review and assessment. *arXiv*  
 783 preprint *arXiv:2312.12148*, 2023.

784

785 Shuai Yang, Liming Jiang, Ziwei Liu, and Chen Change Loy. Pastiche master: Exemplar-based  
 786 high-resolution portrait style transfer. In *CVPR*, 2022.

787

788 Yifan Yang, Jiajun Zhou, Ngai Wong, and Zheng Zhang. Loretta: Low-rank economic tensor-  
 789 train adaptation for ultra-low-parameter fine-tuning of large language models. *arXiv preprint*  
 790 *arXiv:2402.11417*, 2024.

791

792 Juzheng Zhang, Jiacheng You, Ashwinee Panda, and Tom Goldstein. Lori: Reducing cross-task  
 793 interference in multi-task low-rank adaptation. *arXiv preprint arXiv:2504.07448*, 2025.

794

795 Kaichen Zhang, Yifei Shen, Bo Li, and Ziwei Liu. Large multi-modal models can interpret features  
 796 in large multi-modal models. *arXiv preprint arXiv:2411.14982*, 2024a.

797

798 Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng,  
 799 Weizhu Chen, and Tuo Zhao. Adalora: Adaptive budget allocation for parameter-efficient fine-  
 800 tuning. *arXiv preprint arXiv:2303.10512*, 2023.

801

802 Zhi Zhang, Qizhe Zhang, Zijun Gao, Renrui Zhang, Ekaterina Shutova, Shiji Zhou, and Shanghang  
 803 Zhang. Gradient-based parameter selection for efficient fine-tuning. In *Proceedings of the*  
 804 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28566–28577, 2024b.

805

806

807

808

809

810  
811  
812  
813  
814  
815

## Appendix to

# The Chosen Few: Sparse Adaptation for Fine-Tuning Large Models

816 **section A** Sufficient Conditions for STAN Interpretability and Identifiability.  
 817 **section B** More discussion on related methods.  
 818 **section C** Experimental Setting and Runtime Analysis.  
 819 **section D** Additional training loss comparisons.  
 820 **section E** Additional qualitative results for diffusion models.  
 821 **section F** Statistical analysis of sparse latent representations.  
 822 **section G** More visualizations for different STAN injections.  
 823 **section H** Interpreting Disentangled Features in LLM.  
 824 **section I** The Use of Large Language Models.  
 825  
 826  
 827

### A SUFFICIENT CONDITIONS FOR STAN INTERPRETABILITY AND IDENTIFIABILITY

828 **Notation.** Let  $E \in \mathbb{R}^{l \times n}$  and  $D \in \mathbb{R}^{m \times l}$ . For any input  $x \in \mathbb{R}^n$ , define

$$z(x) := Ex, \quad I(x) := T_k(|z(x)|), \quad S(x)_{ii} = \mathbb{1}\{i \in I(x)\}. \quad (7)$$

829 The STAN adaptation operator and output are

$$\mathcal{A}(x) := D S(x) E, \quad \Delta W x = \mathcal{A}(x) x. \quad (8)$$

830 For a standard basis vector  $e_p$  with  $p \in [n]$ , write  $|(Ee_p)_i|$  for the singleton response at latent row  $i$ .

831 **Assumptions used below. They are sufficient, not necessary.**

832 1. Singleton margin:

$$833 \quad i^*(p) := \arg \max_{i \in [l]} |(Ee_p)_i|, \quad \gamma_p := |(Ee_p)_{i^*(p)}| - \max_{j \neq i^*(p)} |(Ee_p)_j| > 0 \quad \text{for all } p \in [n]. \quad (9)$$

834 2. Blockwise identity on singletons:

$$835 \quad D_{\{i^*(p)\}} E_{\{i^*(p)\}} e_p = e_p \quad \text{for all } p \in [n]. \quad (10)$$

836 3. Row orthonormality for identifiability:

$$837 \quad EE^\top = I_l. \quad (11)$$

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 **Theorem 1. Singleton exact recovery implies monosemanticity.** Under Assumptions 1 and 2, for  
 865 every  $p \in [n]$  and every  $\alpha > 0$ ,

$$i^*(p) \in I(\alpha e_p), \quad (12)$$

$$\Delta W(\alpha e_p) = D S(\alpha e_p) E(\alpha e_p) = D_{\{i^*(p)\}} E_{\{i^*(p)\}}(\alpha e_p) = \alpha e_p. \quad (13)$$

866 Moreover, for any  $x$  with  $p \notin \text{supp}(x)$ , one has  $i^*(p) \notin I(x)$ , hence latent  $i^*(p)$  is monosemantic  
 867 for coordinate  $p$ .

871 *Proof. One line.* Since  $\gamma_p > 0$ , Top -  $k$  on  $|E(\alpha e_p)|$  selects  $i^*(p)$ ; substituting the blockwise identity  
 872 yields the claim.

874 **Theorem 2. Identifiability up to permutation and positive scaling.** Assume Assumptions 1 and  
 875 2 for all  $p \in [n]$ , and also assume Assumption 3. If another parameter pair  $(\tilde{D}, \tilde{E})$  satisfies the same  
 876 singleton conditions, then there exist a permutation matrix  $P$  and a positive diagonal matrix  $\Lambda$  such  
 877 that

$$\tilde{E} = \Lambda P E, \quad \tilde{D} = D P^\top \Lambda^{-1}. \quad (14)$$

879 In particular, the set of singleton carriers  $\{i^*(p) : p \in [n]\}$  is unique up to permutation and positive  
 880 rescaling, which means STAN is identifiable modulo permutation and scale.

882 **Remark.** The three assumptions above describe a sufficient regime for interpretability and identifiability.  
 883 They can be checked post hoc by inspecting singleton margins, testing the singleton blockwise  
 884 identity, and measuring  $\|EE^\top - I_l\|_F$ .

## 886 B MORE DISCUSSION ON RELATED METHODS

888 In this section, we differentiate our proposed STAN from other proposed PEFT methods, including  
 889 AdaLoRA (Zhang et al., 2023), and SoRA (Ding et al., 2023).

- 891 • **AdaLoRA (Zhang et al., 2023):**

892 AdaLoRA parameterizes the weight update in an SVD-like form  $\Delta W = P \Lambda Q$ , where  $\Lambda$  contains  
 893 trainable singular values that are adaptively pruned based on sensitivity scores. This allows the  
 894 model to adjust the effective rank of the update matrix during training:

$$W = W^{(0)} + \Delta W = W^{(0)} + P \Lambda Q, \quad (15)$$

897 In contrast, our STAN operates in a high-dimensional latent space and performs adaptation by  
 898 *dynamically* selecting a sparse subset of features for each input using the non-linear topk operator,  
 899 rather than adjusting the overall rank during training. While AdaLoRA seeks a low-rank approxima-  
 900 tion constrained by a parameter budget, STAN focuses on sparse, input-dependent feature selection.  
 901 As a result, STAN enables a more flexible and adaptive mechanism, whereas AdaLoRA offers a  
 902 more static and structured approach.

- 903 • **SoRA (Ding et al., 2023):**

904 SoRA extends the standard LoRA decomposition by introducing a learnable sparse gating vector  
 905  $g \in \mathbb{R}^{r_{\max}}$  between the down- and up-projection matrices. This gating mechanism allows the  
 906 model to adaptively select its effective rank during training. Given an input feature  $x \in \mathbb{R}^q$ , the  
 907 forward pass of a single SoRA block is defined as:

$$z = W_u(g \odot (W_d x)), \quad (16)$$

910 where  $W_d \in \mathbb{R}^{r_{\max} \times q}$  and  $W_u \in \mathbb{R}^{p \times r_{\max}}$  are the projection matrices and  $\odot$  denotes element-  
 911 wise multiplication. The gating is regularized toward sparsity through a proximal-gradient soft-  
 912 thresholding update:

$$g^{t+1} = \mathcal{T}_{\eta_t \lambda}(g^t - \eta_t \nabla_g \mathcal{L}_0(\Delta^t)), \quad \mathcal{T}_\xi(x) = \begin{cases} x - \xi, & x > \xi, \\ 0, & |x| \leq \xi, \\ x + \xi, & x < -\xi, \end{cases} \quad (17)$$

913 where  $\mathcal{L}_0$  is the original task loss,  $\eta_t$  is the learning rate, and  $\lambda$  controls sparsity.

918 SoRA performs static, global pruning of entire rank components. In contrast, STAN adopts a  
 919 dynamic, input-dependent feature selection mechanism using the topk operator, activating different  
 920  $k$ -dimensional subspaces for different inputs without permanently pruning features during training.  
 921 This design offers greater flexibility while maintaining training stability, as evidenced by the results  
 922 shown in Figure 7.

923 • **LS-LoRA** (He et al., 2022):

924 While both methods leverage sparsity, they are fundamentally different in their approach and  
 925 philosophy. LS-LoRA employs a static pruning strategy, using an algorithm like SNIP to remove  
 926 a fixed set of weights once at initialization. Consequently, its sparse pattern is predetermined  
 927 and remains unchanged throughout the entire training process. In contrast, STAN operates on  
 928 the principle of dynamic feature selection. Its sparsity is not static but is instead flexible and  
 929 input-dependent, emerging from the data representations during training. This core difference  
 930 reflects their distinct goals: LS-LoRA focuses on effective compression, whereas STAN is designed  
 931 for conceptual decomposition and representation selection, allowing the model to dynamically  
 932 activate features relevant to a given input.

933 **C EXPERIMENTAL SETTING AND RUNTIME ANALYSIS**

934 In this section, we perform a runtime analysis and detail the experimental settings used throughout our  
 935 evaluations, spanning three task categories: (1) natural language understanding, (2) vision-language  
 936 reasoning, and (3) text-to-image generation.

937 **C.1 RUNTIME ANALYSIS**

938 To provide a comprehensive evaluation of STAN’s practical efficiency, we analyze its performance in  
 939 terms of both training throughput and inference latency.

940 First, we compare STAN against other adaptive sparse fine-tuning methods, namely AdaLoRA and  
 941 SoRA, under the same experimental settings: 3 epochs toy training for SST-2 dataset with 1 A100  
 942 GPU with batch size equals to 32 and sequence length equals to 256 on RoBERTa-base. The results,  
 943 presented in Table 10, demonstrate that STAN achieves superior training throughput compared to both  
 944 baselines. For instance, it processes nearly 50% more tokens per second than AdaLoRA. Crucially,  
 945 this training acceleration is achieved while maintaining an inference latency that is on par with these  
 946 methods, showcasing its superior training efficiency without compromising inference speed.

Method	Inference Latency (ms)	Train Tokens/sec	GPU-hours (3 ep)
STAN	23.16	19815	0.01
AdaLoRA	23.29	13721	0.01
SoRA	25.17	14451	0.015

950 Table 10: Efficiency comparison with adaptive sparse methods AdaLoRA and SoRA. STAN shows  
 951 significantly higher training throughput while maintaining comparable inference latency.

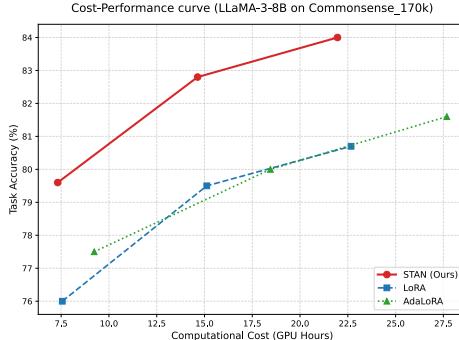
Method	Avg. Inference Latency/sample (s)	Train Tokens/sec	GPU-hours (3 ep)	Peak GPU Memory (GB)
STAN	0.419	2359.42	0.01	39.62
LoRA	0.404	1533.81	0.02	33.53

952 Table 11: Efficiency comparison with LoRA on the LLaVA-1.5-13B model. STAN demonstrates  
 953 superior training efficiency with only a minor trade-off in peak memory.

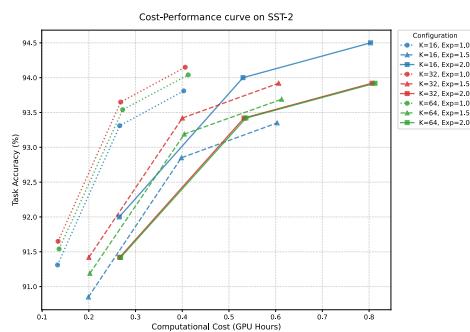
954 To further assess the scalability and practical benefits of our approach, we conduct an additional  
 955 experiment on a larger model, LLaVA-1.5-13B, comparing STAN directly with the widely-used  
 956 LoRA. As detailed in Table 11, STAN maintains a significant advantage in training efficiency over  
 957 LoRA even on this larger scale. Specifically, STAN achieves a higher training throughput and requires

972 approximately half the GPU-hours to complete the same training task. While inference latency is  
 973 comparable to LoRA, STAN incurs only a modest increase in peak GPU memory usage.  
 974

975 To provide a rigorous evaluation of the trade-off between computational cost and task performance,  
 976 we visualize the training trajectories of STAN against baselines.  
 977



989 Figure 5: Cost-Performance curve on  
 990 LLaMA-3-8B.  
 991



989 Figure 6: Robustness analysis across parameter budgets on SST-2.  
 990

991 As illustrated in Figure 5 on the large-scale LLaMA-3-8B benchmark, STAN exhibits superior  
 992 convergence dynamics. It achieves a significantly higher peak accuracy of 84.0% while requiring  
 993 less total wall-clock training time (21.96 GPU hours) compared to LoRA (80.7% accuracy in 22.67  
 994 hours) and AdaLoRA. This empirical evidence highlights that the sparse high-dimensional features in  
 995 STAN facilitate faster learning of task-relevant concepts, effectively offsetting the storage overhead.  
 996

997 Furthermore, Figure 6 investigates the method’s robustness under varying parameter budgets. We  
 998 observe distinct trends in runtime cost: increasing the Expansion Factor (which scales total trainable  
 999 parameters) leads to a proportional increase in GPU hours (shifting curves to the right), while increasing  
 1000  $K$  introduces a minor additional sorting overhead. Despite these variations in computational cost,  
 1001 STAN maintains consistently high performance across different configurations. Notably, efficient  
 1002 settings with lower Expansion Factors (e.g., 1.0) yield optimal accuracy with minimal runtime,  
 1003 confirming that the performance gains stem from the sparse adaptation mechanism itself rather than  
 1004 mere parameter scaling.  
 1005

## 1006 C.2 NATURAL LANGUAGE UNDERSTANDING

1007 **Models and Benchmarks.** We evaluate on five representative tasks from the GLUE benchmark  
 1008 (Wang et al., 2018): MNLI, QNLI, SST-2, CoLA, and QQP. For each task, we follow the  
 1009 official train/test splits. Our experiments are conducted using four pretrained language models:  
 1010 RoBERTa-base, RoBERTa-large (Liu et al., 2019), DeBERTaV3-base, and DeBERTaV3-large (He  
 1011 et al., 2023b).  
 1012

1013 We attribute this compelling performance profile to our design. On large-scale models, the benefits  
 1014 of our dynamic sparsity for accelerating training become more pronounced. During inference, the  
 1015 substantial computational overhead of the base model itself diminishes the relative impact of the  
 1016 small overhead introduced by STAN’s mechanism. These results underscore STAN’s scalability and  
 1017 its practical advantages in large-model scenarios, making it a strong choice for efficient fine-tuning  
 1018 and deployment.  
 1019

1020 **Implementation Details.** During fine-tuning, we apply a weight decay of 0.1 and set the warmup  
 1021 ratio to 0.06. Optimization is performed using the AdamW optimizer with a linear learning rate  
 1022 scheduler. All experiments are conducted with a fixed random seed of 0 for reproducibility. The  
 1023 detailed training hyperparameters for all methods are summarized in Table 12.  
 1024

1025

1026	Method	Model	Epochs	Batch Size	LR	$r / k$	$\alpha$	Lambda	Lambda2
1027	AdaLoRA	RoBERTa-base	15	32	5e-4	2	32	–	–
1028		RoBERTa-large	15	32	5e-4	2	32	–	–
1029		DeBERTaV3-base	15	32	5e-4	2	32	–	–
1030		DeBERTaV3-large	15	32	5e-4	2	32	–	–
1031	SORA	RoBERTa-base	15	32	8e-4	8	16	10	3e-4
1032		RoBERTa-large	15	32	8e-4	8	16	10	3e-4
1033		DeBERTaV3-base	15	32	8e-4	8	16	10	3e-4
1034		DeBERTaV3-large	15	32	8e-4	8	16	10	3e-4
1035	LoRA	RoBERTa-base	15	32	5e-4	8	16	–	–
1036		RoBERTa-large	15	32	5e-4	8	16	–	–
1037		DeBERTaV3-base	15	32	5e-4	8	16	–	–
1038		DeBERTaV3-large	15	32	5e-4	8	16	–	–
1039	STAN (Ours)	RoBERTa-base	15	32	5e-4	8	–	–	–
1040		RoBERTa-large	15	32	5e-4	8	–	–	–
1041		DeBERTaV3-base	15	32	5e-4	8	–	–	–
1042		DeBERTaV3-large	15	32	5e-4	8	–	–	–

Table 12: Summary of training hyperparameters used for each method and backbone model. Parameters not applicable to a given method are indicated with –.

### C.3 VISION-LANGUAGE REASONING

**Models and Benchmarks.** We conduct our experiments using three versions of the LLaVA (Liu et al., 2023a) model: LLaVA\_1.5\_7B, LLaVA\_1.5\_13B, and LLaVA\_1.6\_Vicuna\_13B. We fine-tune these models on two benchmarks, GQA (Hudson & Manning, 2019) and ScienceQA (Lu et al., 2022), for each task we follow the official train/test splits provided.

**Implementation Details.** For fine-tuning the LLaVA models, we use the LLaVA\_1.5\_13B checkpoint equipped with the CLIP-ViT-L/14-336 vision encoder and an MLP-based projector. Training is conducted using DeepSpeed ZeRO-3 for efficient large-scale optimization. Detailed training hyperparameters for all LLaVA models and methods are provided in Table 13.

1058	Method	Model	Epochs	Batch Size	LR	$r / k$	$\alpha$	Num Latents	Warmup Ratio
1059	LoRA	LLaVA_1.5_7B	15	16	2e-4	128	256	–	0.03
1060		LLaVA_1.5_13B	15	16	2e-4	128	256	–	0.03
1061		LLaVA_1.6_Vicuna-13B	15	16	2e-4	128	256	–	0.03
1062	STAN (Ours)	LLaVA_1.5_7B	15	16	2e-4	128	–	512	0.03
1063		LLaVA_1.5_13B	15	16	2e-4	128	–	512	0.03
1064		LLaVA_1.6_Vicuna-13B	15	16	2e-4	128	–	512	0.03

Table 13: Summary of training hyperparameters for each method and backbone model. Parameters not applicable to a given method are indicated with –.

### C.4 TEXT-TO-IMAGE GENERATION

**Models and Benchmarks.** Experiments are conducted using the publicly available Stable Diffusion 3 (SD3) (Rombach et al., 2022) released by Stability AI, trained on a composite dataset comprising images from WikiArt (Saleh & Elgammal, 2015) and the DualStyleGAN dataset (Yang et al., 2022). The WikiArt dataset contains a broad collection of real-world artworks spanning diverse styles and historical periods, while DualStyleGAN provides synthetically generated images with fine-grained, diverse, controllable stylistic attributes.

**Implementation Details.** The training hyperparameters corresponding to all diffusion models and methods are summarized in Table 14.

1080	Method	Model	Resolution	Steps	Batch Size	Grad Accum.	LR	$r / k$	$\alpha$	Num Latents
1081	LoRA	SD3	1024	3000	1	1	4e-4	32	1	-
1082	STAN (Ours)	SD3	1024	3000	1	1	4e-4	96	-	512

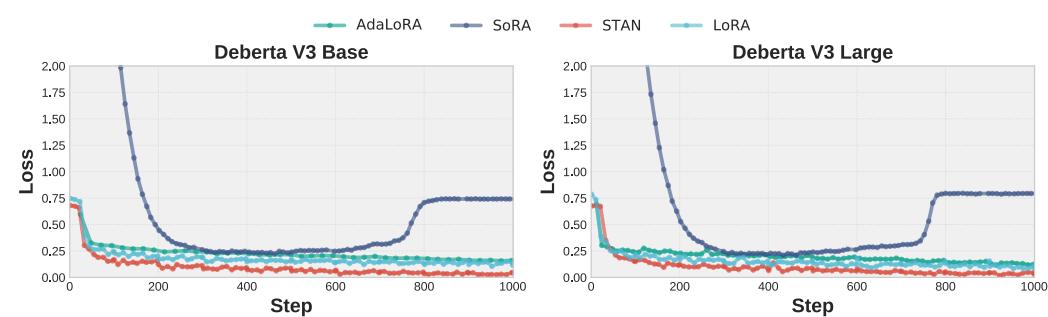
1084  
1085 Table 14: Summary of training hyperparameters for each method and backbone model. Parameters  
1086 not applicable to a given method are indicated with -.

## 1088 D ADDITIONAL TRAINING LOSS COMPARISONS

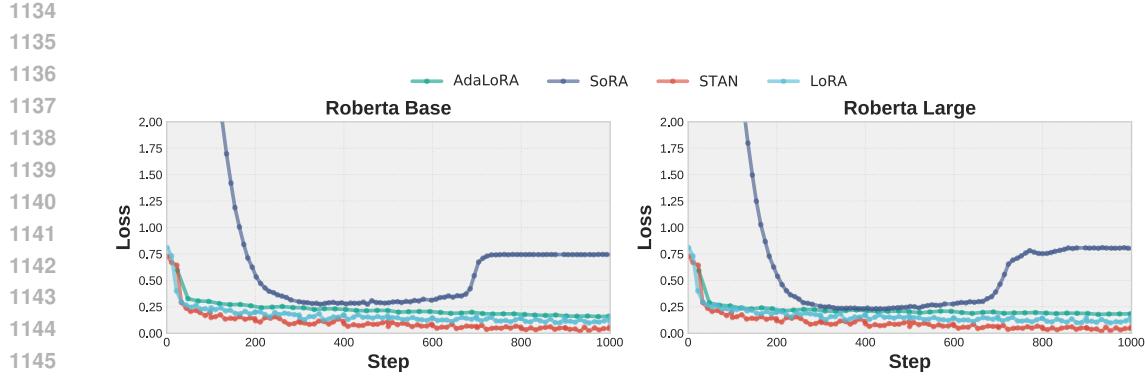
1090 Beyond performance, the stability of the fine-tuning process is a critical indicator of a PEFT method's  
1091 robustness and reliability. As previously highlighted in Table I, although methods such as SoRA (Ding  
1092 et al., 2023) incorporate sparsity, their practical utility is often limited by unstable training dynamics.

1093 In this section, we present additional loss curves for comparison. Across all demonstrations, STAN  
1094 consistently exhibits stable convergence, whereas other methods display either significant oscillations  
1095 or fail to converge effectively. Notably, in some cases, the loss for SoRA and LoRA approaches  
1096 zero, suggesting potential overfitting during the fine-tuning process. This indicates a limitation  
1097 in representational capacity, highlighting the challenges of capturing task-relevant features under  
1098 constrained adaptation schemes.

1099 To further investigate this, The following Figures present training loss curves on the different datasets  
1100 for STAN, compared against three other PEFT methods across different base models. Across all  
1101 model architectures, STAN exhibits remarkably stable and smooth convergence. The training loss  
1102 decreases consistently, without significant oscillations, and generally reaches or surpasses the final  
1103 loss values achieved by LoRA and AdaLoRA. In contrast, the loss curves for SoRA display notable  
1104 volatility and often plateau at higher values or even diverge – corroborating earlier observations of  
1105 its instability. This erratic behavior during training likely contributes to SoRA's inconsistent and  
1106 sometimes suboptimal task performance. By comparison, the stable convergence of STAN not only  
1107 strengthens confidence in its training dynamics but also underpins the consistent, high performance  
1108 reported in prior sections. These findings suggest that STAN's approach to sparse adaptation is  
1109 inherently more robust, making it a more reliable and effective choice for parameter-efficient fine-  
1110 tuning.



1122 Figure 7: Training loss comparison on SST-2 using DeBERTaV3-base/large (He et al., 2023b). STAN  
1123 consistently achieves the lowest loss, demonstrating greater stability compared to other methods.



1164 Figure 9: Loss comparison on QNLI using RoBERTa-base and RoBERTa-large (Liu et al., 2019).  
1165 STAN achieves the lowest and most stable loss curve. LoRA exhibits more oscillations and signs  
1166 of overfitting toward the end of training. In contrast, both AdaLoRA and SoRA show unstable  
1167 convergence and fail to reach optimal performance.  
1168  
1169  
1170  
1171  
1172      AdaLoRA      SoRA      STAN      LoRA  
1173 **Debertа V3 Base**      Loss  
1174      Step  
1175  
1176  
1177  
1178 **Debertа V3 Large**      Loss  
1179      Step  
1180  
1181  
1182

1183 Figure 10: Loss comparison on QNLI using DeBERTaV3-base and DeBERTaV3-large (He et al.,  
1184 2023b). STAN achieves the lowest and most stable loss trajectory. LoRA exhibits greater oscillation  
1185 during training, while AdaLoRA shows signs of overfitting early on. SoRA remains unstable  
1186 throughout and also displays indications of overfitting.  
1187

1188 E ADDITIONAL QUALITATIVE RESULTS FOR DIFFUSION MODELS  
11891190 In this section, we present additional images generated by SD3, fine-tuned using different methods  
1191 and conditioned on different prompts.

1192

1193

1194 **STYLE** : abstract expressionism1195 **CONTENT** : ladies with watermelons

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210



(a) STAN



(b) LoRA



(c) None

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



(a) STAN



(b) LoRA



(c) None

1242  
1243  
1244  
1245  
1246  
1247

STYLE : baroque  
CONTENT : women in robes



(a) STAN



(b) LoRA



(c) None

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

STYLE : cartoon

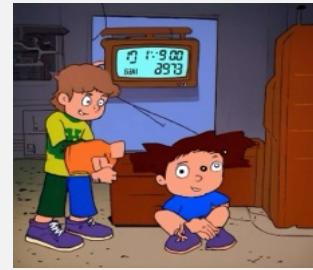
CONTENT : two kids building a time machine in a garage



(a) STAN



(b) LoRA



(c) None

1277

1278

1279

1280

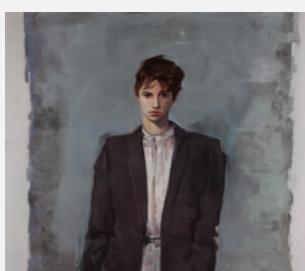
1281

1282

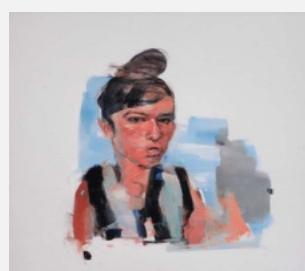
1283

STYLE : ecole de paris

CONTENT : portrait



(a) STAN



(b) LoRA



(c) None

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296

1297

1298

1299

STYLE : expressionism

CONTENT : snowy ground



(a) STAN



(b) LoRA



(c) None

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

STYLE : fantasy

CONTENT : a castle carved into a cliff glowing with runes



(a) STAN



(b) LoRA



(c) None

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

STYLE : illustration

CONTENT : a rainy day with two friends playing chess



(a) STAN



(b) LoRA



(c) None

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350

1351

1352

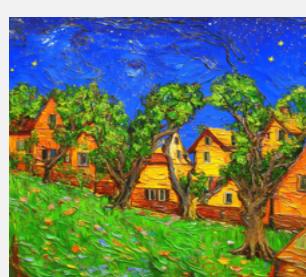
1353 **STYLE** : impasto

1354 **CONTENT** : a village with twisted trees and stars overhead

1355



(a) STAN



(b) LoRA



(c) None

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372 **STYLE** : impressionism

1373 **CONTENT** : woman in the garden



(a) STAN



(b) LoRA



(c) None

1383

1384

1385

1386

1387

1388

1389

1390 **STYLE** : naive art primitivism

1391 **CONTENT** : blue eyed flowers



(a) STAN



(b) LoRA



(c) None

1401

1402

1403

1404

1405

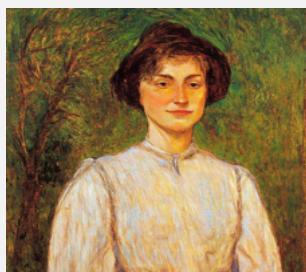
1406

1407

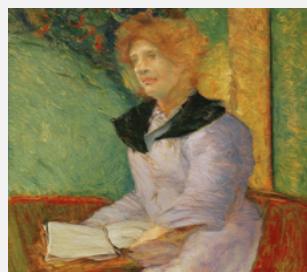
STYLE : neo impressionism

CONTENT : a women

1409



(a) STAN



(b) LoRA



(c) None

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

STYLE : neoclassicism

CONTENT : two women



(a) STAN



(b) LoRA



(c) None

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

STYLE : post impressionism

CONTENT : a portrait of a man inside a frame



(a) STAN



(b) LoRA



(c) None

1455

1456

1457

1458

1459

1460

1461 **STYLE** : pre raphaelite brotherhood

1462 **CONTENT** : three female figures{dancing and playing

1463



(a) STAN



(b) LoRA



(c) None

1464

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

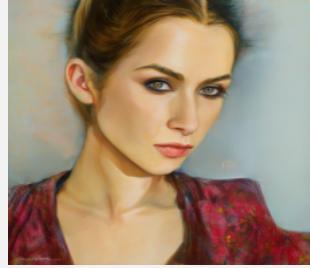
1477

1478

1479

1480

1481



(a) STAN



(b) LoRA



(c) None

1491

1492

1493

1494

1495

1496

1497

1498

**STYLE** : rococo

**CONTENT** : a man



(a) STAN



(b) LoRA



(c) None

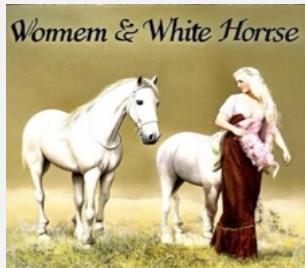
1509

1510

1511

1512  
1513  
1514  
1515  
1516  
1517

STYLE : romanticism  
CONTENT : a woman and a white horse



(a) STAN



(b) LoRA

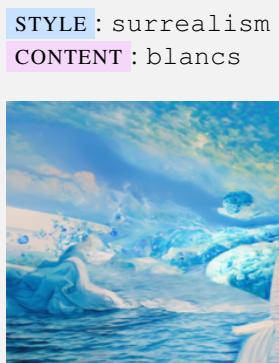


(c) None

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



(a) STAN



(b) LoRA



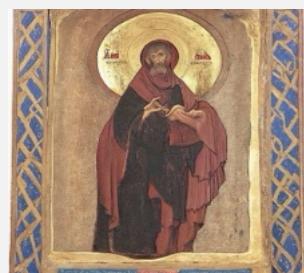
(c) None

1545  
1546  
1547  
1548  
1549

1550  
1551  
1552  
1553  
1554



(a) STAN



(b) LoRA



(c) None

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

STYLE : ukiyoe

CONTENT : Mount Fuji behind blooming cherry trees



(a) STAN



(b) LoRA



(c) None

## 1620 F STATISTICAL ANALYSIS OF SPARSE LATENT REPRESENTATIONS

1622 In this section, we conduct a statistical analysis of the sparse latent representations learned by STAN,  
 1623 as visualized in Figure 3. Our goal is to quantify the behavior and specialization of individual latent  
 1624 units by measuring their activation patterns across different input samples and semantic classes.  
 1625 Specifically, we evaluate three statistical properties for each latent unit: the frequency of activation  
 1626 (x-axis), the mean activation value (y-axis), and the label entropy (color bar), which reflects the class  
 1627 specificity of each latent.

1628 Let  $\{a_i\}_{i=1}^N$  denote the activation values of a given latent unit across  $N$  input samples. We consider a  
 1629 latent unit to be *active* on sample  $i$  if  $a_i > 0$ . Let  $\mathcal{I} = \{i \mid a_i > 0\}$  represent the index set of active  
 1630 samples, and let  $N_+ = |\mathcal{I}|$  be the number of activations. Each input sample is associated with a  
 1631 ground truth class label  $y_i \in \mathcal{C}$ , where  $\mathcal{C}$  denotes the set of all class labels.

1633 **(1) Mean Activation Frequency.** We compute the mean frequency with which the latent unit is  
 1634 activated across the dataset as:

$$1635 \text{MeanFreq} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[a_i > 0] = \frac{N_+}{N}. \quad (18)$$

1638 This metric reflects how often the latent contributes non-zero activation to the representation.

1640 **(2) Mean Activation Value.** For the subset of samples where the latent is active, we compute the  
 1641 average activation magnitude:

$$1642 \text{MeanAct} = \frac{1}{N_+} \sum_{i \in \mathcal{I}} a_i. \quad (19)$$

1645 This value captures the typical strength of activation, conditioned on the latent being active.

1647 **(3) Label Entropy.** To assess how class-specific the latent activation is, we compute a label-  
 1648 weighted activation distribution. First, we define the activation proportion  $p_c$  for each class  $c \in \mathcal{C}$   
 1649 as:

$$1650 p_c = \frac{\sum_{i:y_i=c} a_i, \mathbf{1}[a_i > 0]}{\sum_{j=1}^N a_j, \mathbf{1}[a_j > 0]}. \quad (20)$$

1652 Then, we compute the entropy of this distribution:

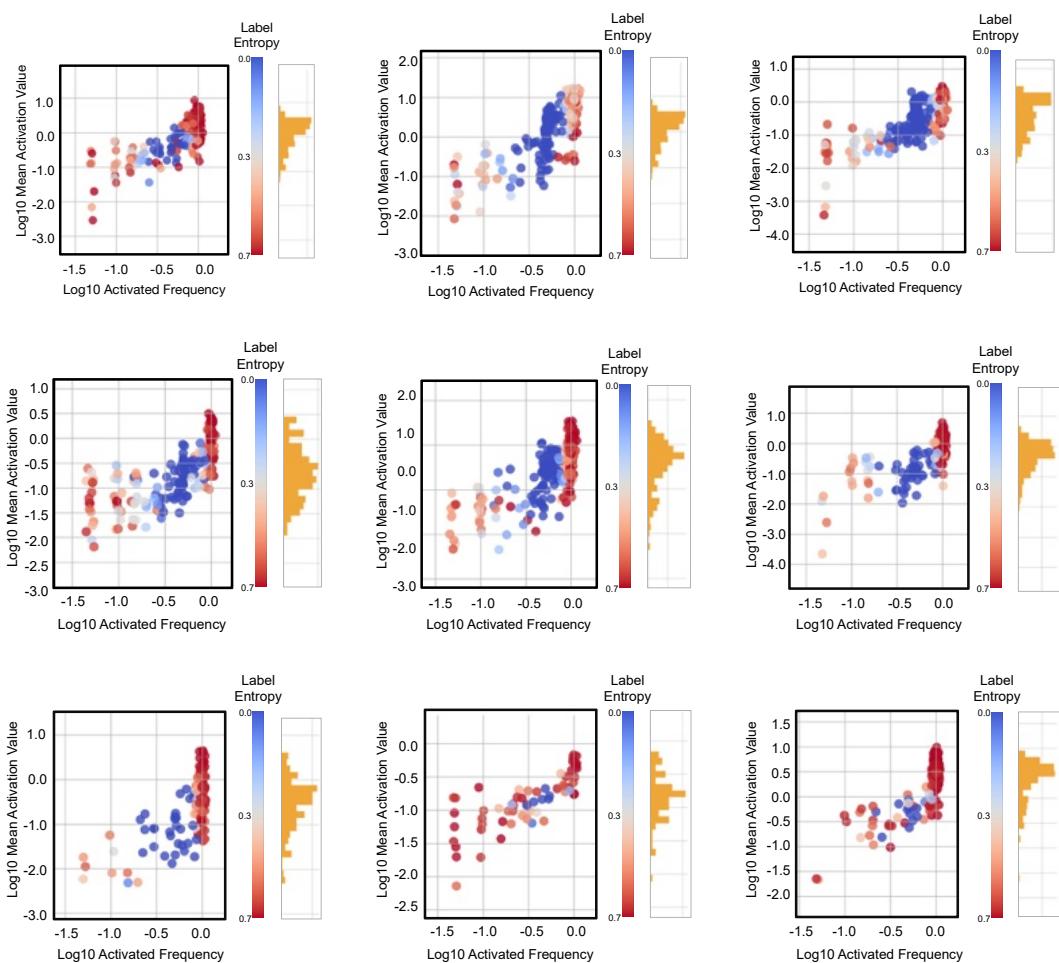
$$1654 \text{Entropy} = - \sum_{c \in \mathcal{C}} p_c \log p_c. \quad (21)$$

1656 Lower entropy values indicate that the latent is primarily activated by a narrow set of classes, implying  
 1657 class specialization. Conversely, higher entropy suggests that the latent is shared across multiple  
 1658 classes, capturing more general or abstract features.

1659 These three metrics together allow us to characterize both the usage pattern and semantic specificity  
 1660 of each latent dimension in STAN. Additional examples and visualizations of these statistics are  
 1661 provided in Appendix G.

## 1674 G MORE VISUALIZATIONS FOR DIFFERENT STAN INJECTIONS

1675  
 1676 In this section, we present additional visualizations consistent with those shown in Figure 3, con-  
 1677 structed using the statistical analysis formulations described in Appendix F. These supplementary  
 1678 results provide further evidence that both sparsity and broad representational capacity are consistently  
 1679 exhibited across different injection layers within the model. This consistency indicates that the  
 1680 behavior of STAN is not limited to a specific configuration, but instead reflects a generalizable  
 1681 property of the method. The observed uniformity and stability across injections support the reliability  
 1682 and robustness of STAN.

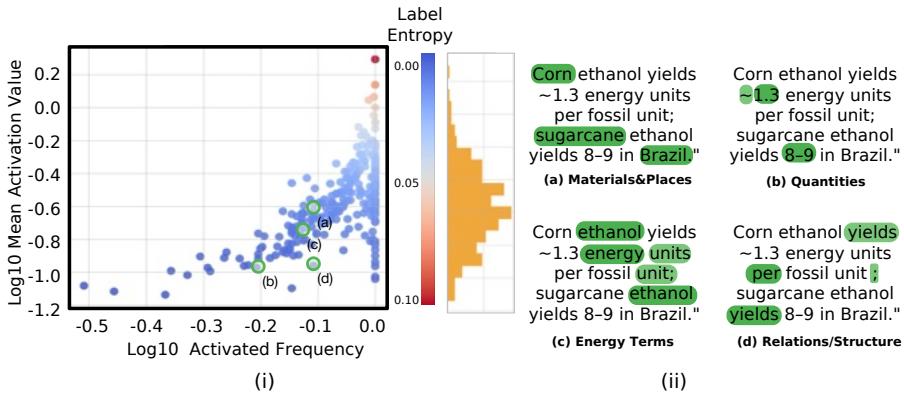


1714 **Figure 32: Visualization examples of sparse representations across different STAN injection**  
 1715 **layers.** Each subplot presents the distribution of middle-layer feature activations, with both axes  
 1716 scaled using  $\log_{10}$  to improve visual clarity. The color bar denotes label entropy, where lower values  
 1717 indicate class-specific activations and higher values reflect class-general behavior, as defined in  
 1718 Appendix F. To the right of each plot, a yellow bar chart illustrates the average activation magnitude  
 1719 for each unit, providing additional insight into the strength and distribution of sparsely activated  
 1720 features.

1721  
 1722  
 1723  
 1724  
 1725  
 1726  
 1727

1728 **H INTERPRETING DISENTANGLED FEATURES IN LLM**  
1729  
1730  
1731

1732 In this section, we present additional interpretability experiment on language model utilizing the  
1733 similar setting shown in Figure 3 and Appendix F



1746 Figure 33: Qualitative demonstration of the sparse representations learned by STAN. (i) Distribution of  
1747 sparsely activated intermediate features across four distinct styles. (ii) Highlight of the corresponding  
1748 text concepts: (a) Materials & Places, (b) Quantities, (c) Energy Terms and (d) Relations / Structure.  
1749

1750 Part (i) illustrates activation frequency statistics of sparsely activated intermediate features when  
1751 model is fed with different texts. The y-axis represents  $\log_{10}$  mean activation value, and the x-axis indicates  
1752  $\log_{10}$  activation frequency. Each point is colored according to its label entropy, which  
1753 reflects the degree of concept specificity associated with that neuron: lower entropy values indicate  
1754 specialization (i.e., the neuron is primarily activated by one concept), while higher entropy suggests  
1755 activation by a broader mix of styles. Details on the formulation of this analysis can be found in  
1756 Appendix F. Part (ii) presents corresponding visual examples, illustrating that the texts activating a  
1757 given neuron are consistently aligned with its associated concept. The results reveal that different  
1758 concepts tend to activate distinct, often non-overlapping, subsets of sparse latent features. A small  
1759 number of shared neurons appear to capture common generative priors, while the majority remain  
1760 concept-specific. It demonstrates the expressive capacity of subspace combinations in the sparse  
1761 latent space.

1762 **I THE USE OF LARGE LANGUAGE MODELS**  
1763

1764 In line with the ICLR policy, we disclose the use of Large Language Models during the preparation of  
1765 this manuscript. Our use of these tools was strictly limited to assistance with language and formatting.  
1766 Specifically, we employed an LLM for proofreading, correcting grammatical errors, and improving  
1767 the clarity and readability of sentences. The LLM had no role in the core scientific aspects of this  
1768 work, including research ideation, methodological design, experimental analysis, or the generation of  
1769 any results or conclusions. All intellectual contributions and the core content of this paper are solely  
1770 the work of the authors.

1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781