

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 SOLAR: COMMUNICATION-EFFICIENT MODEL ADAP- TATION VIA SUBSPACE-ORIENTED REPARAMETRIZA- TION

006 **Anonymous authors**

007 Paper under double-blind review

011 ABSTRACT

013 Parameter-efficient fine-tuning (PEFT) methods, such as LoRA, enable scalable
014 adaptation of foundation models by injecting low-rank adapters. However, their
015 communication and storage costs remain a major bottleneck in resource-constrained
016 settings. We propose **SOLAR** (Subspace-Oriented Latent Adapter Reparameter-
017 ization), a post-training compression framework that substantially reduces the
018 communication cost (i.e., the number of parameters to transmit or store) of PEFT
019 adapters. SOLAR expresses each PEFT update as a linear combination of basis vec-
020 tors formed from the foundation model’s singular vectors with controlled random
021 perturbations. By exploiting the subspace similarity (the alignment of principal
022 directions) between the foundation model and task-specific fine-tuned updates,
023 SOLAR decouples the adapter size from PEFT structure and ensures compact yet
024 expressive representations. It is model-agnostic and compatible with existing PEFT
025 methods, including LoRA, AdaLoRA, and other adapter modules. We theoretically
026 establish a bound on the reconstruction error. Experiments on language and vision
027 tasks using LLaMA, GPT, and ViT models demonstrate that SOLAR preserves
028 task performance while significantly reducing model representation sizes, offering
029 an effective and communication-efficient solution for deployment in distributed
030 systems and edge devices.

032 1 INTRODUCTION

033 Foundation models—large-scale pretrained transformer architectures—have catalyzed substantial
034 progress across natural language processing, computer vision, and a range of other domains. However,
035 adapting these models to downstream tasks remains resource-intensive. Full fine-tuning, which
036 updates all model parameters, demands considerable computational, memory, and storage resources
037 Houldsby et al. (2019). Parameter-Efficient Fine-Tuning (PEFT) techniques address this challenge by
038 freezing the backbone and updating only a small set of task-specific parameters. For example, adapter
039 modules insert compact trainable layers into each network block Houldsby et al. (2019); prefix-tuning
040 optimizes a continuous prompt of only $\sim 0.1\%$ of the model’s parameters Li & Liang (2021); and
041 Low-Rank Adaptation (LoRA) injects low-rank update matrices into each layer Hu et al. (2021).
042 These methods achieve performance comparable to fully fine-tuned models while updating less than
043 1% of the model’s parameters.

044 Despite these parameter savings, the cumulative communication and storage costs of PEFT modules
045 remain a critical bottleneck in many real-world scenarios, particularly as foundation models continue
046 to scale Wolf et al. (2020). In distributed scenarios (e.g., federated learning), these adapters must be
047 communicated and stored across multiple devices or nodes, leading to significant overhead Wolf et al.
048 (2020). Communication and storage overhead increase with the number of PEFT modules, as many
049 fine-tuned adapters are saved and frequently transmitted or synchronized, thus turning millions of
050 adapter parameters into a major bottleneck, particularly in bandwidth-limited or memory-constrained
051 environments such as edge devices or federated learning systems Gao & Zhang (2024); Wang et al.
052 (2025). The resulting communication and storage costs (i.e., the number of adapter parameters that
053 must be transmitted and stored) can lead to slower training, increased energy consumption, and
reduced scalability, highlighting the need for more efficient adapter compression techniques.

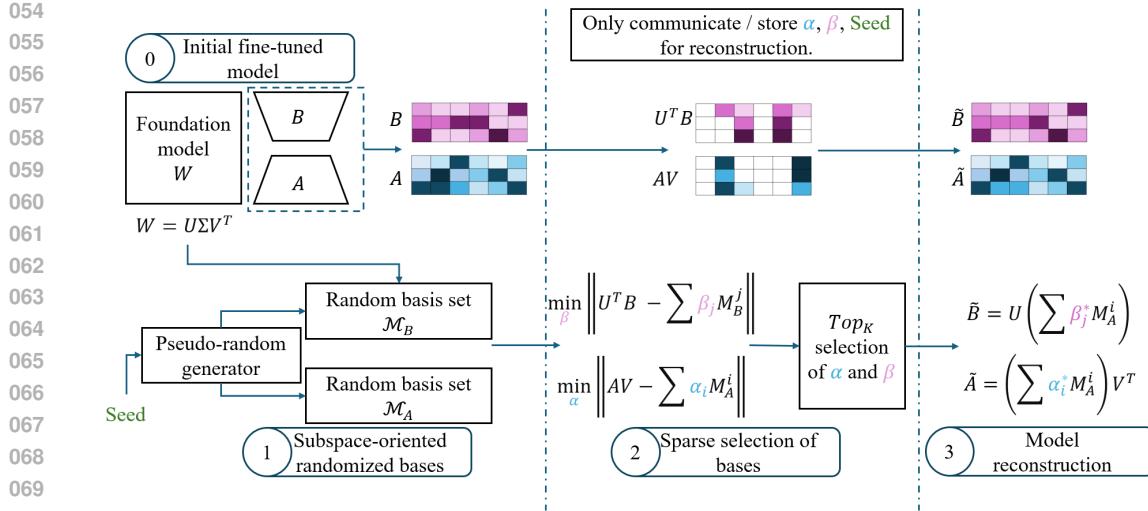


Figure 1: Overview of SOLAR. Given fine-tuned adapters (A, B), SOLAR projects them onto structured subspaces derived from the pretrained model’s SVD. A seeded pseudo-random generator (seeded with a known value) deterministically creates the basis matrices. Top- k coefficients α and β are selected under a budget to reconstruct \tilde{A} and \tilde{B} , while the bases are never stored or transmitted. Only the coefficients α, β , and the seed need to be communicated or stored.

To address this, several methods decouple tunable parameters from adapter rank and model dimensions: NOLA Koohpayegani et al. (2024) expresses LoRA’s matrices as linear combinations of random basis matrices, training only the coefficients; VeRA Kopiczko et al. (2023) uses shared frozen random vectors with small learned scaling vectors; and SVFT Lingam et al. (2024) constructs a basis from singular vectors of pretrained weights and learns a sparse combination during fine-tuning. However, random bases not aligned with the model or task may reduce representational efficiency, and methods such as Kopiczko et al. (2023); Lingam et al. (2024); Koohpayegani et al. (2024) are not post-hoc, as they modify the training process and cannot compress adapters already trained—creating a need for a flexible, training-free compression utility.

In this paper, we propose SOLAR (Subspace-Oriented Latent Adapter Reparameterization), a novel post-training compression method for PEFT adapters. SOLAR exploits the empirical structure of adapter updates by reparameterizing them as linear combinations of structured, randomized basis matrices. It is model-agnostic and applicable post-training without modifying the fine-tuning process. The main contributions of this work are as follows:

- We leverage the observed subspace similarity between the foundation model’s weights (W) and the task-specific update (ΔW) to create a more compact and efficient adapter representation. By expressing ΔW as a sparse combination of basis vectors, our method effectively decouples the adapter’s final size from the model’s architecture.
- We develop a three-step framework for post-hoc adapter compression that involves: 1) constructing a basis pool of size N by perturbing the foundation model’s singular vectors with random noise, 2) performing a sparse selection of the most significant basis vectors to meet a budget k , and 3) reconstructing the adapter using only the selected coefficients and a single random seed.
- We provide a formal theoretical analysis that bounds the reconstruction error. Our proof decomposes the total error into the original training error and a controllable compression error, which can be minimized by tuning SOLAR’s hyperparameters (N and k).
- We demonstrate through extensive experiments that SOLAR reduces adapter sizes by up to 98% while preserving the performance of the original LoRA adapters. Our results show competitive accuracy across a wide range of vision and language tasks using ViT, GPT-2, and LLaMA models.

2 PROPOSED METHOD: SOLAR

We propose a *post-training* compression strategy that serves as a modular add-on for compressing PEFT-based updates. It introduces no training overhead and is compatible with LoRA Hu et al. (2021),

108 QLoRA Dettmers et al. (2023), Compacter Karimi Mahabadi et al. (2021), and NOLA Koohpayegani
 109 et al. (2024), operating post-hoc by taking the final trained adapter matrices as input. SOLAR applies
 110 to OFT Qiu et al. (2023) and variants Liu et al. (2023), compressing $\Delta W = (R - I)W$ via its
 111 SVD-based subspace without altering the orthogonal parameterization. By exploiting the low-rank
 112 structure of updates, SOLAR significantly reduces communication and storage costs in distributed or
 113 resource-limited settings.

114

115

2.1 PROBLEM FORMULATION

116

117 Transformer-based models parameterize attention and MLP layers using full-rank weight matrices
 118 $W \in \mathbb{R}^{m \times n}$. Recent PEFT methods, such as LoRA Hu et al. (2021), decompose the task-specific
 119 update ΔW as $\Delta W = BA$, where $A \in \mathbb{R}^{r \times n}$, $B \in \mathbb{R}^{m \times r}$, and $r \ll \min(m, n)$. This reduces the
 120 trainable parameters from mn to $r(m + n)$, yielding a compression ratio of $\frac{mn}{r(m+n)}$. While effective,
 121 LoRA’s fixed-rank formulation limits its flexibility. Alternatives, such as NOLA Koohpayegani
 122 et al. (2024), leverage random projections to approximate ΔW , but often require large basis sets to
 123 sufficiently capture the relevant directions. To address this challenge and enhance compression further,
 124 we formulate the problem as minimizing the approximation loss between ΔW and its compressed
 125 counterpart $\tilde{\Delta W}$ subject to a strict communication (or storage) budget:

126

$$\min_{\tilde{\Delta W}} \|\Delta W - \tilde{\Delta W}\|_F^2, \quad \text{s.t. } \|\tilde{\Delta W}\|_0 \leq k, \quad (1)$$

127

128 where $\|\cdot\|_F$ denotes the Frobenius norm, and $\|\cdot\|_0$ represents the number of non-zero elements (i.e.,
 129 $\|X\|_0 \triangleq \sum_{i=1}^m \sum_{j=1}^n \mathbb{I}\{X_{ij} \neq 0\}$). The parameter k specifies the total budget.

130

131 Building on the LoRA formulation, we approximate the individual factors A and B , aiming to find
 132 compressed counterparts \tilde{A}, \tilde{B} such that:

133

$$\min_{\tilde{A}, \tilde{B}} \|BA - \tilde{B}\tilde{A}\|_F^2, \quad \text{s.t. } \|\tilde{A}\|_0 \leq k_A, \quad \|\tilde{B}\|_0 \leq k_B, \quad k_A + k_B = k, \quad (2)$$

134

135 where k_A and k_B represent budgets for \tilde{A} and \tilde{B} , respectively. This problem is challenging: counting
 136 the number of nonzero elements is non-convex, sparse element selection is combinatorial, and
 137 excessive sparsity may degrade accuracy. Achieving high compression without task performance loss
 thus requires careful subspace design and adaptive optimization.

138

139

140

2.2 METHOD: SUBSPACE-ORIENTED RANDOMIZED BASIS, SPARSE SELECTION, AND RECONSTRUCTION

141

142 To solve (2), we propose SOLAR. A key insight motivating our approach is that ΔW predominantly
 143 resides in the subspace spanned by W , particularly in LoRA-based fine-tuning, where constraining the
 144 rank $r \ll \min(m, n)$ forces ΔW to concentrate its variation along specific directions of W Hu et al.
 145 (2021). This alignment (i.e., the overlap in the principal directions of W and ΔW) has been observed
 146 empirically and explained theoretically via neural tangent kernel (NTK) theory Jacot et al. (2018);
 147 Malladi et al. (2023); Seleznova et al. (2023). The left- and right-singular alignments are measured
 148 as $\|U_W^\top U_{\Delta W}\|_F^2$ and $\|V_W^\top V_{\Delta W}\|_F^2$, where U and V contain the left and right singular vectors from
 149 the SVD of each matrix Hu et al. (2021). Under this perspective, the model’s response to updates
 150 is well-approximated by a first-order expansion: $f(\xi; W + \Delta W) \approx f(\xi; W) + \langle \nabla f(\xi; W), \Delta W \rangle$,
 151 where f is the model, ξ is input data, and $\nabla_W f(\xi; W)$ denotes the gradient of the foundation model’s
 152 output. This implies that ΔW lies in a low-curvature (and hence low-dimensional) subspace defined
 153 by W ’s parameter space (see Section 3.4 for empirical evidence). Thus, projecting ΔW into the
 154 subspace of W enables an efficient and compact representation that can be sparsified with minimal
 155 information loss.

156

157

158 Building on these insights, we design a three-stage compression framework (Figure 1). First, we
 159 construct a randomized basis set aligned with the foundation model (Section 2.2.1). Next, we select
 160 a sparse set of bases to approximate the projected update (Section 2.2.2). We then reconstruct the
 161 update using a budget-aware combination of selected components (Section 2.2.3).

162

163

2.2.1 STEP 1: SUBSPACE-ORIENTED RANDOMIZED BASIS SET

164

165

166 We construct a basis set from the foundation model’s parameter space via SVD of the model weight,
 167 $W = U\Sigma V^T$, where $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ are orthonormal, and $\Sigma \in \mathbb{R}^{m \times n}$ is diagonal. This

162 decomposition enables a basis naturally aligned with the directions of task-specific updates ΔW .
 163 Unlike methods such as NOLA Koohpayegani et al. (2024) relying on unstructured random bases,
 164 our foundation-aligned directions allow a more compact representation of ΔW .

165 To enrich the expressive power of this subspace, we construct randomized basis matrices by perturbing
 166 slices of the singular vectors:

$$168 \quad \mathcal{M}_A = \left\{ M_A^{(i)} = V[:, \mathcal{I}_i] + \epsilon_i \right\}_{i=1}^{N_A}, \quad \mathcal{M}_B = \left\{ M_B^{(j)} = U[:, \mathcal{J}_j] + \epsilon_j \right\}_{j=1}^{N_B}, \quad (3)$$

170 where \mathcal{I}_i and \mathcal{J}_j are randomly sampled index sets, N_A, N_B are the number of basis candidates for A
 171 and B , respectively, and ϵ_i, ϵ_j are random matrices with each entry drawn i.i.d. from $\mathcal{N}(0, 1)$. These
 172 basis sets form a flexible pool of candidates for approximation.

174 2.2.2 STEP 2: SPARSE SELECTION OF BASES

176 To enable more compact approximations, the LoRA update $\Delta W = BA$ is first projected into the
 177 subspace of W . Given the singular value decomposition $W = U\Sigma V^T$, this projection is defined as
 178 $\Delta W_{\text{Proj}} = U^T \Delta W V = (U^T B)(AV) = B_{\text{Proj}} A_{\text{Proj}}$, where $A_{\text{Proj}} = AV$ and $B_{\text{Proj}} = U^T B$ represent
 179 the update components expressed in the basis of W . This transformation retains all information
 180 when W is full-rank, and is particularly effective when ΔW is already aligned with the foundation
 181 subspace, a property commonly observed in LoRA-based fine-tuning. Under this projection, the
 182 update becomes $\Delta W = U \Delta W_{\text{Proj}} V^T$. This approach leverages the inherent alignment between W
 183 and ΔW , enabling more efficient approximations with fewer basis elements than methods such as
 184 NOLA, which rely on unstructured random projections. Specifically, we approximate the projected
 185 LoRA factors AV and $U^T B$ using sparse linear combinations of the basis matrices:

$$186 \quad \min_{\alpha} \left\| AV - \sum_{i=1}^{N_A} \alpha_i M_A^{(i)} \right\|_F^2, \text{ s.t. } \|\alpha\|_0 \leq k_A, \quad \min_{\beta} \left\| U^T B - \sum_{j=1}^{N_B} \beta_j M_B^{(j)} \right\|_F^2, \text{ s.t. } \|\beta\|_0 \leq k_B. \quad (4)$$

190 A two-step strategy is employed to solve these NP-hard problems efficiently. The first step computes
 191 the unconstrained least squares solution to obtain coefficients α^* and β^* . The second step applies
 192 hard thresholding to retain only the top k entries by magnitude based on the budgets k_A and k_B .

194 2.2.3 STEP 3: BUDGET-AWARE RECONSTRUCTION

196 The approximated model update is then reconstructed using the selected top k bases, resulting in \tilde{A}
 197 and \tilde{B} for A and B , respectively:

$$198 \quad A \approx \left(\sum_{i \in S_A} \alpha_i^* M_A^{(i)} \right) V^T, \quad B \approx U \left(\sum_{j \in S_B} \beta_j^* M_B^{(j)} \right), \quad (5)$$

200 where S_A and S_B are the selected top k index sets. Because the update reconstruction is performed
 201 within the subspace defined by W , this step ensures strong alignment with task-relevant direc-
 202 tions. The reconstruction balances accuracy and compression, with the sparsity budgets k_A and k_B
 203 controlling the number of active basis.

204 **Adaptive Compression.** SOLAR enables flexible allocation of sparsity budgets k_A and k_B , adapting
 205 to system constraints such as memory, storage, or bandwidth. This allows deployment on resource-
 206 constrained devices, with adapter size dynamically adjustable post-training. For instance, a server
 207 can send a compact adapter to low-memory clients and a richer version to more capable devices.

209 2.3 THEORETICAL ANALYSIS OF RECONSTRUCTION ERROR

211 We assume that (A1) the model is initialized with spectral initialization; (A2) the optimal update is
 212 low-rank; (A3) the change in the model’s weights from fine-tuning is well-behaved according to the
 213 generation process in Zhang et al. (2025a); and (A4) the singular values of the projected update matrix
 214 exhibit Fast Spectrum Decay. These assumptions are well-established and frequently utilized in the
 215 literature for convergence analyses, as in previous works, such as Zhang et al. (2025a); Martinsson &
 Tropp (2020).

216 **Theorem 1 [SOLAR Reconstruction Error Bound]** Let ΔW^* be the optimal low-rank adapter,
 217 ΔW be the adapter learned via fine-tuning, and $\Delta \tilde{W}$ be the adapter reconstructed by SOLAR.
 218 Under assumptions (A1)–(A4), the expected total error is bounded by $\mathbb{E} \left[\|\Delta \tilde{W} - \Delta W^*\|_F \right] \leq$
 219 $C_1 + C_2$, where C_1 captures the fine-tuning error (depending on learning rate, training steps,
 220 and spectrum of ΔW^* ; see Appendix A), and $C_2 = \sqrt{1 + \frac{r_A}{N_A - r_A - 1}} \left(\sum_{t > r_A} \sigma_t^2(\Delta W) \right)^{\frac{1}{2}} +$
 221 $\sqrt{1 + \frac{r_B}{N_B - r_B - 1}} \left(\sum_{t > r_B} \sigma_t^2(\Delta W) \right)^{\frac{1}{2}} + \left(\sum_{t > r_A} \sigma_t^2(\Delta W) \right)^{\frac{1}{2}}$, where $\sigma_t(\Delta W)$ is the t -th singular
 222 value of the fine-tuned update ΔW , and r_A, r_B denote the effective ranks after moving to the random
 223 basis space. The SOLAR reconstruction error has two parts: the fine-tuning error (C_1) and the
 224 compression error (C_2). The compression error decreases with larger basis pools (N_A, N_B) and
 225 higher sparsity budget (k). Details are in Appendix A.
 226

229 3 EXPERIMENTS

230 We evaluate SOLAR through extensive experiments in three domains: 1) image classification with
 231 ViT-B/L in few-shot and full-data settings (Section 3.1); 2) instruction tuning on LLaMA-3 models
 232 using Alpaca and MMLU (Section 3.2); and 3) language generation with GPT-2 on E2E NLG
 233 (Section 3.3). Across all settings, SOLAR matches LoRA and NOLA in accuracy while reducing
 234 adapter size by up to 98%, offering a lightweight representation for model adaptation.
 235

236 3.1 SOLAR ON VISION TRANSFORMERS

237 We conduct few-shot image classification experiments using ViT-B and ViT-L Dosovitskiy et al.
 238 (2020) foundation models, initialized with either supervised or self-supervised He et al. (2022).
 239

240 **Experimental Setup.** We compare SOLAR against LoRA Hu et al. (2021) and NOLA Koohpayegani
 241 et al. (2024). Experiments are conducted on ViT-Base (ViT-B) and ViT-Large (ViT-L) architectures.
 242 Supervised ViT models pretrained on ImageNet-21k Deng et al. (2009) are obtained from Google’s
 243 official releases via the Hugging Face repository Wolf et al. (2020); Research (2025), and MAE
 244 models pretrained on ImageNet-1K are sourced from the Timm library Wightman (2025). All
 245 experiments run on a single NVIDIA RTX 4090 GPU using PyTorch Paszke (2019) and HuggingFace
 246 libraries. In SOLAR, the compressed representation consists of (i) a random seed to regenerate
 247 the basis vectors, (ii) an encoded list of selected basis indices, and (iii) their coefficients. Reported
 248 trainable parameters include both projection coefficients and overhead (i.e., seed and index encoding).
 249 The MLP classifier head is dataset-specific and excluded from the parameter count unless noted.
 250

251 **Evaluation Benchmarks.** We fine-tune on standard image classification datasets: CIFAR-10
 252 Krizhevsky et al. (2009), CIFAR-100 Krizhevsky et al. (2009), Food-101 Bossard et al. (2014),
 253 Tiny-ImageNet Le & Yang (2015), ImageNet-1K Deng et al. (2009), Oxford Pets Parkhi et al. (2012),
 254 SUN397 Xiao et al. (2010), and CUB-200-2011 Welinder et al. (2010).
 255

256 **Comparison Methods.** We compare SOLAR with several baselines: Full Fine-Tuning (Full-FT),
 257 LoRA Hu et al. (2021), and NOLA Koohpayegani et al. (2024). In Full-FT, all backbone parameters
 258 are updated. For LoRA, we apply low-rank adapters to the attention Query projection matrices,
 259 with a rank of 4 for ViT-B and either 1 or 4 for ViT-L. For NOLA, following Koohpayegani et al.
 260 (2024), adapters are inserted into MLP layers using 1000 random basis vectors for each of the A
 261 and B matrices. All models are trained with cross-entropy loss. For full-data settings, we train 5
 262 epochs with batch size 128; for few-shot settings (10 samples per class), 25 epochs with batch size
 263 16, emphasizing low-data efficiency relevant to real-world and distributed scenarios. To account for
 264 variance from limited data, we sample four training splits per dataset and report mean top-1 accuracy
 265 on the test split (or validation for ImageNet-1k). Experiments are repeated with different random
 266 seeds, and learning rates are tuned per dataset and model. Additional details are in the appendix.
 267

268 **Results and Performance Analysis.** We evaluate SOLAR on various vision benchmarks using foun-
 269 dation models, with results in Table 1. In the tables, configurations are denoted as $\text{SOLAR}_{\text{method}(N \rightarrow k)}$,
 indicating that SOLAR is applied to a NOLA or LoRA model trained with rank r , using N bases per
 270 matrix ($N = N_A = N_B$) and selecting the top- k bases by significance, where N and k are given in
 271 thousands. SOLAR consistently achieves competitive top-1 accuracy in few-shot (10 samples per

270 Table 1: Top-1 classification accuracy (%) of ViT-B and ViT-L on benchmark datasets under two
 271 settings: (1) few-shot (10 samples/class, 25 epochs) and (2) full-data (5 epochs). Results report
 272 mean \pm std over 5 runs. SOLAR is applied with configuration $\text{method}(N \rightarrow k)$, where N and k are in
 273 thousands.

Model	Method	# Param	CIFAR-10		CIFAR-100		Food-101		T-ImageNet	
			10	Full	10	Full	10	Full	10	Full
ViT-B	Full-FT	86M	91.1 \pm .8	94.6 \pm .5	78.2 \pm .7	87.7 \pm .3	65.8 \pm .9	85.2 \pm .4	78.1 \pm .10	85.4 \pm .6
	LoRA ($r=4$)	74K	92.3 \pm .6	98.3 \pm .2	81.8 \pm .8	90.3 \pm .4	72.4 \pm .7	87.6 \pm .3	77.9 \pm .9	88.8 \pm .4
	NOLA	48K	92.2 \pm .6	94.7 \pm .5	81.3 \pm .8	86.6 \pm .4	72.6 \pm .5	85.9 \pm .2	78.4 \pm .7	82.8 \pm .5
	SOLAR $_{r=4(4 \rightarrow 1.6)}$	41K	92.3 \pm .7	98.3 \pm .4	<u>81.5</u> \pm .7	<u>89.8</u> \pm .2	71.8 \pm .6	<u>87.0</u> \pm .5	77.9 \pm .8	<u>87.9</u> \pm .4
	SOLAR $_{\text{NOLA}(4 \rightarrow 1.2)}$	32K	92.1 \pm .7	94.5 \pm .3	81.1 \pm .6	85.4 \pm .3	72.5 \pm .6	85.4 \pm .3	<u>78.3</u> \pm .8	82.3 \pm .5
ViT-L	Full-FT	303M	90.2 \pm .9	94.1 \pm .6	86.2 \pm .7	87.7 \pm .5	73.9 \pm .8	85.5 \pm .4	80.8 \pm .11	89.2 \pm .6
	LoRA ($r=4$)	197K	97.1 \pm .5	98.7 \pm .1	88.1 \pm .7	<u>92.4</u> \pm .3	81.8 \pm .7	89.8 \pm .2	84.4 \pm .8	91.8 \pm .5
	LoRA ($r=2$)	98K	<u>96.6</u> \pm .4	98.7 \pm .1	88.0 \pm .6	92.9 \pm .3	<u>82.1</u> \pm .7	90.0 \pm .2	83.8 \pm .7	90.4 \pm .3
	NOLA	96K	96.0 \pm .8	97.4 \pm .6	87.8 \pm 1.0	89.3 \pm .5	82.5 \pm .8	86.7 \pm .4	84.3 \pm .9	86.7 \pm .6
	SOLAR $_{r=4(4 \rightarrow 1.6)}$	82K	97.0 \pm .5	<u>98.5</u> \pm .3	<u>87.9</u> \pm .8	91.4 \pm .4	76.8 \pm .7	<u>87.1</u> \pm .4	78.7 \pm .7	88.6 \pm .5
	SOLAR $_{r=2(1 \rightarrow 0.3)}$	50K	96.1 \pm .8	98.2 \pm .4	87.4 \pm .9	90.0 \pm .5	77.0 \pm .8	86.8 \pm .6	76.4 \pm .9	87.6 \pm .6
	SOLAR $_{\text{NOLA}(4 \rightarrow 1.2)}$	<u>64K</u>	95.8 \pm .9	97.0 \pm .4	87.7 \pm .8	89.3 \pm .4	<u>82.1</u> \pm .7	86.6 \pm .3	<u>84.1</u> \pm .8	86.4 \pm .6

282 Table 2: Additional evaluation on vision datasets using ViT-B. The table shows bit-level representation
 283 footprint (32-bit baseline) and top-1 accuracy. All models are trained for 10 epochs.

Method	Byte Footprint	Oxford Pets	SUN397	CUB-200	ImageNet-1K
LoRA ($r=1$)	74KB	93.0 \pm .3	74.3 \pm .2	84.7 \pm .2	81.5 \pm .4
NOLA	48KB	90.4 \pm .5	61.7 \pm .4	79.4 \pm .4	77.4 \pm .3
SOLAR $_{r=1(2 \rightarrow 0.2)}$	8KB (89% \downarrow)	<u>92.6</u> \pm .4	<u>73.9</u> \pm .2	84.2 \pm .3	81.3 \pm .2

297 Table 3: Effect of quantization on
 298 SOLAR $_{r=4(4 \rightarrow 1.6)}$ performance. ViT-L-
 299 MAE fine-tuned on CIFAR-10.

Method	Quant.	Accuracy	Byte Footprint
SOLAR	32-bit	86.7 \pm .3	319KB
	16-bit	86.5 \pm .3	166KB
	8-bit	85.9 \pm .4	89KB
	4-bit	84.8 \pm .6	50KB

297 Table 4: Effect of rank and adapter placement in
 298 SOLAR $_{r=4(4 \rightarrow 1)}$. Accuracy (%) on CIFAR-100
 299 using ViT-B.

Rank	Q	K	V	QV	QKV
1	87.0	85.5	86.6	88.3	90.1
2	87.5	85.7	87.4	88.6	90.5
4	87.8	86.1	87.5	89.0	90.6
8	88.1	86.0	87.4	89.1	90.7
16	87.9	86.0	87.1	89.0	90.6

309 class) and full-data settings while requiring far fewer trainable parameters than LoRA and NOLA.
 310 On ViT-B and ViT-L, SOLAR matches LoRA’s performance using up to 74% fewer parameters. For
 311 instance, applied to a LoRA ($r = 2$), bases $N_A = N_B = 4000$, and $\text{top}_k = 1600$, SOLAR reduces
 312 fine-tuned parameters from 98K to 25K while maintaining comparable accuracy.

313 Beyond parameter reduction, SOLAR improves storage efficiency. Table 2 reports mean and standard
 314 deviation over 5 runs on four additional datasets using ViT-B, quantifying the bit-level footprint
 315 assuming 32-bit precision during training. We apply 8-bit quantization to SOLAR after top_k parameter
 316 selection. While LoRA ($r = 1$) requires 74KB of adapter parameters, SOLAR reduces this to 8KB
 317 (89% reduction). These extreme compressions incur only minor accuracy drops, showing SOLAR
 318 enables fine-grained control of model size to meet strict constraints and offers a flexible tradeoff
 319 between footprint and performance.

320 In addition to reducing parameter and storage footprints, SOLAR remains highly robust under
 321 quantization. As shown in Table 3, reducing coefficient precision from 32-bit to 4-bit incurs less than
 322 a 2% accuracy drop on ViT-L-MAE (CIFAR-10, 10-shot). We further evaluate the effect of adapter
 323 rank and placement (Table 4), observing that performance improves with rank up to 8 (with higher
 324 ranks requiring more time to converge), and that the Query (Q) projection yields the highest gains.

324
 325 Table 5: Model representation efficiency for LLaMA models. SOLAR compresses LoRA adapter
 326 updates across various model sizes. For the 13B model, all methods use 4-bit quantization, making
 327 the LoRA baseline equivalent to QLoRA.

Model	LLaMA-3.2 1B			LLaMA-2 13B (4-bit)		
	LoRA $r=8$	NOLA 1000 bases	SOLAR $r = 8(4 \rightarrow 1.2)$	LoRA $r=1$	NOLA 1000 bases	SOLAR $r = 1(1 \rightarrow 0.3)$
# Params	852K	64K	81K (90% ↓)	819K	140K	51K (94% ↓)
Val Loss	1.51	1.87	1.52	1.05	1.29	1.05
MMLU Acc	30.1	25.9	28.3	54.5	51.8	54.5

335
 336 Table 6: Performance and parameter efficiency on E2E NLG using GPT-2 Small and Medium. All
 337 methods use rank-4 adapters applied to the Query and Value projections.

Method	GPT-2 Small		GPT-2 Medium	
	MET	# Params	MET	# Params
Full-FT	28.4	124M	46.2	355M
LoRA ($r=4$)	29.7	147K	47.2	393K
NOLA	29.1	48K	46.8	350K
SOLAR ($r=4, 1 \rightarrow 0.3$)	29.7	15K (90% ↓)	46.4	30K (92% ↓)
SOLAR ($r=1, 0.1 \rightarrow 0.1$)	26.1	4K (97% ↓)	44.8	9K (98% ↓)

3.2 SOLAR ON LLaMA

348
 349 **Experimental Setup.** We apply SOLAR to LLaMA-3 models of size 1B–13B. All models are fine-
 350 tuned using adapters in the query and value projections across all transformer layers. For the 1B model,
 351 we use LoRA with rank 8; for the 31B model, we use LoRA with rank 1. To reduce GPU memory
 352 usage for large-scale models, we quantize the 13B model using 4-bit NF4 quantization through the
 353 `BitsAndBytes` library Dettmers et al. (2021); Dettmers (2025). Further implementation details
 354 and hardware configurations are provided in the Appendix.

355
 356 **Evaluation Benchmarks.** All models are fine-tuned on the Stanford Alpaca Taori et al. (2023)
 357 dataset for instruction-following and evaluated on its validation loss. We also assess generalization to
 358 out-of-distribution tasks using the MMLU benchmark Hendrycks et al. (2020).

359
 360 **Comparison Methods.** We compare SOLAR with PEFT baselines, including LoRA Hu et al. (2021)
 361 and NOLA Koohpayegani et al. (2024). LoRA uses rank $r = 8$ for LLaMA-3 1B and $r = 1$
 362 for the 13B model. NOLA follows its original configuration, with 1000 random basis vectors per
 363 matrix Koohpayegani et al. (2024). For the 13B model, we apply 4-bit quantization to all methods
 364 (LoRA, NOLA, and SOLAR). The reported trainable parameters include learned coefficients and
 365 overhead for basis indexing. All experiments use gradient checkpointing, and learning rates are tuned
 366 separately per model and method to ensure a fair comparison.

367
 368 **Results and Performance Analysis.** Table 5 reports results across model sizes. SOLAR matches
 369 LoRA in Alpaca validation loss and MMLU Hendrycks et al. (2020) accuracy while reducing trainable
 370 adapter parameters by up to 94%. For example, on LLaMA-3.2 13B, SOLAR cuts the adapter size
 371 from 819K to 51K without accuracy loss.

3.3 SOLAR ON GPT-2

372
 373 **Experimental Setup.** We evaluate our method on GPT-2 Radford et al. (2019) base and medium
 374 models fine-tuned on the E2E NLG dataset Novikova et al. (2017) using LoRA. The models are
 375 trained for 5 epochs using a batch size of 8 and a learning rate of 0.1. LoRA is applied to the
 376 self-attention Query and Value projection, with a rank of $r = 4$. After training, we apply SOLAR to
 377 compress the LoRA adapter updates.

378
 379 **Evaluation Benchmarks.** We use the E2E NLG dataset to evaluate generative quality. Generated
 380 outputs are assessed using METEOR Banerjee & Lavie (2005) metric. We report LoRA, NOLA, and
 381 SOLAR performance.

378 **Results and Performance Analysis.** Table 6 summarizes results on the E2E NLG dataset using
 379 GPT-2 Small and Medium models. SOLAR achieves competitive METEOR scores compared to
 380 LoRA and NOLA, while substantially reducing adapter size. On GPT-2 Medium, SOLAR reduces
 381 adapter representation size from 393K (LoRA) to 30K parameters with minimal performance loss.
 382 Applied to rank-1 LoRA, it achieves a 98% reduction, demonstrating strong compression capability.
 383

384 3.4 DISCUSSION AND ANALYSIS ON SOLAR PERFORMANCE AND EFFICIENCY

385 **Subspace Analysis.** We analyze the sub-
 386 space similarity between the foundation model’s
 387 weights W and the LoRA update ΔW with rank
 388 $r = 4$ (see Figure 2). Let $W = U_W \Sigma_W V_W^\top$ and
 389 $\Delta W = U_{\Delta W} \Sigma_{\Delta W} V_{\Delta W}^\top$ denote their SVDs.
 390 To quantify subspace alignment, we define
 391 the similarity function as $\phi(W, \Delta W, i, j) =$
 392 $\psi(U_W^{(i)}, U_{\Delta W}^{(j)}) = \|U_W^{(i)^\top} U_{\Delta W}^{(j)}\|_F^2$, where $U_W^{(i)}$
 393 and $U_{\Delta W}^{(j)}$ are the matrices formed by taking
 394 the i and j left singular vectors of W and
 395 ΔW , respectively. This normalized Frobenius
 396 inner product measures how much of the j -
 397 dimensional subspace of ΔW lies within the
 398 i -dimensional subspace of W , reaching its max-
 399 imum when perfectly aligned. Figure 2 shows
 400 that the fine-tuned model emphasizes directions already present in the foundation model, supporting
 401 prior observations that LoRA updates lie in low-dimensional, structured subspaces Hu et al. (2021);
 402 Farhadzadeh et al. (2025); Zhang et al. (2025b). This suggests leveraging existing directions is
 403 more effective than relying purely on random ones: LoRA implicitly aligns with them, and SOLAR
 404 exploits this alignment in its basis pool, explaining its performance advantage over NOLA.
 405

406 Effect of Basis Pool Size and Communication

407 **Budget on Performance.** To evaluate SOLAR’s
 408 trade-off between representation size and per-
 409 formance, we analyze the effect of varying the
 410 basis pool size and the number of selected top $_k$
 411 components on representation accuracy. Experi-
 412 ments are conducted on a ViT-Base model fine-
 413 tuned using LoRA with rank 4, followed by SO-
 414 LAR compression. Each LoRA matrix A and
 415 B requires $4 \times 768 = 3072$ parameters. We
 416 observe that increasing k improves SOLAR’s
 417 expressiveness and accuracy. Moreover, a larger
 418 basis pool enhances performance by increasing
 419 the likelihood of capturing directions aligned
 420 with the fine-tuned model subspace. As shown
 421 in Figure 3, even with fixed k , larger pools yield higher accuracy by enabling more precise reconstruc-
 422 tion of target directions. SOLAR thus achieves performance comparable to LoRA with significantly
 423 fewer parameters. This trade-off confirms Theorem 1: increasing the basis pool N or sparsity k
 424 reduces the compression error C_2 .

425 **SOLAR Overhead and Runtime Efficiency.** As a
 426 post-training method, SOLAR introduces negligible
 427 runtime overhead and does not interfere with fine-
 428 tuning. For instance, fine-tuning LLaMA-3.2 1B with
 429 LoRA on Tiny-ImageNet took 2081 seconds, while
 430 SOLAR, including random basis generation, convex
 431 least-squares solving, and top $_k$ selection, took only
 432 15 seconds (under 0.72% of training time). These
 433 operations are computationally lightweight, as shown
 434 in Table 7, confirming SOLAR’s practical efficiency.

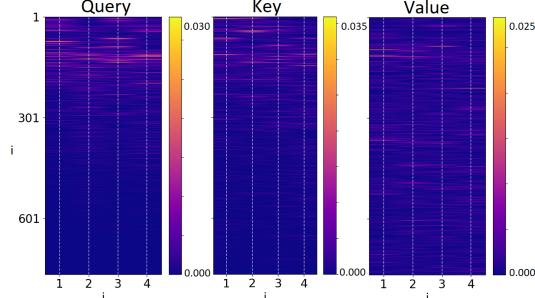


Figure 2: Subspace similarity between the W and ΔW matrices (Q, K, V) from the first layer of the ViT-B model using LoRA with rank $r = 4$.

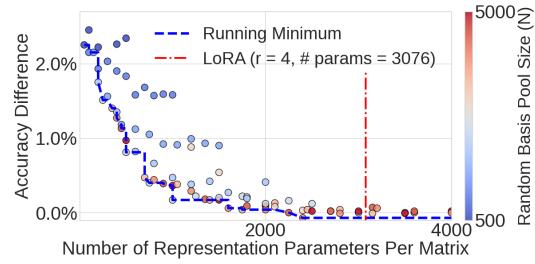


Figure 3: Representation Performance vs. Representation Cost: On ViT-B ($r = 4$), SOLAR demon-
 strates a trade-off between parameter count and
 performance, achieving strong performance with
 far fewer parameters than LoRA.

Table 7: Runtime Overhead: LoRA (10
 epochs) vs. SOLAR post-training on ViT-B
 across vision datasets. Times in seconds.

Dataset	LoRA	SOLAR	Overhead (%)
CIFAR-10	1176	14	1.19
CIFAR-100	1165	14	1.20
Food-101	3480	67	1.92
Tiny-ImageNet	2081	15	0.72
ImageNet-1K	56634	155	0.27

432 **Limitations and Future Work.** As a post-hoc method, SOLAR’s performance is limited by the base
 433 adapter, and its hyperparameters (N and k) may need per-task tuning to optimize the compression-
 434 accuracy trade-off. While it shows strong results on vision and language tasks, its effectiveness on
 435 other modalities (audio, time series, or multimodal data) remains untested. Future work will extend
 436 SOLAR to these areas and evaluate its performance in other environments.
 437

438 4 BACKGROUND AND RELATED WORKS

439 **Transformers in NLP and Vision.** Transformers Vaswani et al. (2017), are now the standard in
 440 NLP for modeling long-range dependencies via self-attention Raiaan et al. (2024). Models such as
 441 LLaMA Touvron et al. (2023), BERT Devlin et al. (2019), and GPT Radford et al. (2018) build on
 442 this structure to achieve strong results across diverse benchmarks. In vision, ViT Dosovitskiy et al.
 443 (2020) treats image patches as tokens, making Transformers a unifying backbone across modalities.
 444

445 **Parameter-Efficient Fine-Tuning (PEFT).** As transformers scale, task-specific fine-tuning becomes
 446 computationally intensive. PEFT methods mitigate this by updating only a subset of parameters.
 447 LoRA Hu et al. (2021) introduces trainable low-rank matrices per layer, typically modifying <1%
 448 of weights, while NOLA Koohpayegani et al. (2024) re-parameterizes these as linear combinations
 449 of random bases, decoupling parameters from rank and architecture. Yet PEFT gains often fall
 450 short in deployment, especially on edge, mobile, and federated settings with communication and
 451 storage bottlenecks. Adapting GPT-2 (117M) on-device may still require gigabytes of transfer and
 452 petaflop-scale computation per round Wang et al. (2025), with updates taking seconds to transmit and
 453 hours to process on low-power hardware (e.g., Jetson TX2).
 454

455 **Challenges of PEFT.** As models grow, adapter overhead scales rapidly. Even modest adapters (e.g.,
 456 7M parameters for a 7B model at rank 16) accumulate significant costs across users, tasks, or training
 457 rounds Xu et al. (2023b). A 1% adapter for LLaMA-2 70B adds 700M parameters; for GPT-3 (350B),
 458 3.5B—tens of gigabytes in FP32. Such costs are infeasible in personalized or federated settings,
 459 where hundreds of adapters may be exchanged or stored per user Zhang et al. (2024). While PEFT
 460 leverages the low intrinsic dimensionality of task adaptation Hu et al. (2021), deployment remains
 461 inefficient. It has been shown that BERT fine-tuning on MRPC Dolan & Brockett (2005) requires only
 462 1,861 degrees of freedom out of 110M, highlighting redundancy in full-rank updates Aghajanyan et al.
 463 (2020). Yet even small adapters impose substantial overhead on massive models Xu et al. (2023a);
 464 Lialin et al. (2023). Hence, the true bottleneck is adapter size, not fine-tuning efficiency Jie et al.
 465 (2023), motivating flexible post-training compression to reduce footprint without altering training.
 466

467 **PEFT Compression Techniques.** To mitigate PEFT costs, pruning Han et al. (2024); Ilhan et al.
 468 (2024) and quantization Chen et al. (2024); Hubara et al. (2021) have been explored. These reduce
 469 model size but require careful tuning or retraining, are less effective under severe bandwidth limits,
 470 and are mainly optimized for full-model compression, limiting applicability to adapters. Adapter
 471 updates are highly redundant and lie in low-dimensional subspaces Hu et al. (2021); Yadav et al.
 472 (2023); Wu et al. (2024), motivating post-training compression. Methods like ComPEFT Yadav
 473 et al. (2023), BitDelta Liu et al. (2024), Delta-CoMe Ping et al. (2024), and DeltaZip Yao et al.
 474 (2025) compress adapter weights after fine-tuning but rely on heuristics, task-specific tuning, or
 475 training integration, reducing flexibility. Other approaches alter fine-tuning itself: VeRA Kopiczko
 476 et al. (2023) employs a shared random basis, SVFT Lingam et al. (2024) learns sparse coefficients
 477 for an SVD-based basis, and EigenLoRAX Kaushik et al. (2025) builds a PCA basis from many
 478 pre-trained adapters. In contrast, SOLAR is a post-hoc, training-free utility that compresses any
 479 adapter, providing a complementary plug-and-play solution.
 480

481 5 CONCLUSION

482 Adapter-based fine-tuning methods such as LoRA significantly reduce the cost of adapting large
 483 models. However, in distributed and on-device settings, communication and storage overheads remain
 484 a major bottleneck. To address this, we introduce SOLAR, a lightweight post-training compression
 485 method that reparameterizes adapter updates as sparse combinations of structured basis vectors
 486 aligned with the foundation model’s latent subspace. SOLAR substantially reduces adapter size and
 487 transmission cost without altering the training process or model architecture.
 488

486 REFERENCES
487

488 Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. Intrinsic dimensionality explains the
489 effectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255*, 2020.

490 Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved
491 correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic*
492 *evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.

493 Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101–mining discriminative compo-
494 nents with random forests. In *Computer vision–ECCV 2014: 13th European conference, zurich,*
495 *Switzerland, September 6–12, 2014, proceedings, part VI 13*, pp. 446–461. Springer, 2014.

496 Mengzhao Chen, Wenqi Shao, Peng Xu, Jiahao Wang, Peng Gao, Kaipeng Zhang, and Ping Luo.
497 Efficientqat: Efficient quantization-aware training for large language models. *arXiv preprint*
498 *arXiv:2407.11062*, 2024.

499 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale
500 hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
501 pp. 248–255. Ieee, 2009.

502 Tim Dettmers. Bitsandbytes: 8-bit optimizers and quantization. <https://github.com/TimDettmers/bitsandbytes>, 2025. Accessed: 15-May-2025.

503 Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise
504 quantization. *arXiv preprint arXiv:2110.02861*, 2021.

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

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

511 Bill Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In
512 *Third international workshop on paraphrasing (IWP2005)*, 2005.

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

517 Farzad Farhadzadeh, Debasmit Das, Shubhankar Borse, and Fatih Porikli. Lora-x: Bridging founda-
518 tion models with training-free cross-model adaptation. *arXiv preprint arXiv:2501.16559*, 2025.

519 Chao Gao and Sai Qian Zhang. Dlora: Distributed parameter-efficient fine-tuning solution for large
520 language model. *arXiv preprint arXiv:2404.05182*, 2024.

521 Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness:
522 Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):
523 217–288, 2011.

524 Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. Parameter-efficient fine-tuning
525 for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*, 2024.

526 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked
527 autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*
528 *vision and pattern recognition*, pp. 16000–16009, 2022.

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

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

543

544 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 545 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint*
 546 *arXiv:2106.09685*, 2021.

547

548 Itay Hubara, Yury Nahshan, Yair Hanani, Ron Banner, and Daniel Soudry. Accurate post training
 549 quantization with small calibration sets. In *International Conference on Machine Learning*, pp.
 549 4466–4475. PMLR, 2021.

550

551 Fatih Ilhan, Gong Su, Selim Furkan Tekin, Tiansheng Huang, Sihao Hu, and Ling Liu. Resource-
 552 efficient transformer pruning for finetuning of large models. In *Proceedings of the IEEE/CVF*
 553 *Conference on Computer Vision and Pattern Recognition*, pp. 16206–16215, 2024.

554

555 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and
 556 generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

557

558 Shibo Jie, Haoqing Wang, and Zhi-Hong Deng. Revisiting the parameter efficiency of adapters
 559 from the perspective of precision redundancy. In *Proceedings of the IEEE/CVF International*
 559 *Conference on Computer Vision*, pp. 17217–17226, 2023.

560

561 Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank
 562 hypercomplex adapter layers. *Advances in Neural Information Processing Systems*, 34:1022–1035,
 2021.

563

564 Prakhar Kaushik, Ankit Vaidya, Shravan Chaudhari, and Alan Yuille. Eigenlorax: Recycling adapters
 565 to find principal subspaces for resource-efficient adaptation and inference. In *Proceedings of the*
 566 *Computer Vision and Pattern Recognition Conference*, pp. 649–659, 2025.

567

568 Soroush Abbasi Koohpayegani, KL Navaneet, Parsa Nooralinejad, Soheil Kolouri, and Hamed
 569 Pirsavash. Nola: Compressing lora using linear combination of random basis. *ICLR 2024*, 2024.

570

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

572

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

574

575 Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

576

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

578

579 Vladislav Lialin, Vijeta Deshpande, and Anna Rumshisky. Scaling down to scale up: A guide to
 579 parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.15647*, 2023.

580

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

584

585 James Liu, Guangxuan Xiao, Kai Li, Jason D Lee, Song Han, Tri Dao, and Tianle Cai. Bitdelta: Your
 586 fine-tune may only be worth one bit. *Advances in Neural Information Processing Systems*, 37:
 587 13579–13600, 2024.

588

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

591

592 Sadhika Malladi, Alexander Wettig, Dingli Yu, Danqi Chen, and Sanjeev Arora. A kernel-based
 593 view of language model fine-tuning. In *International Conference on Machine Learning*, pp.
 23610–23641. PMLR, 2023.

594 Per-Gunnar Martinsson and Joel A Tropp. Randomized numerical linear algebra: Foundations and
 595 algorithms. *Acta Numerica*, 29:403–572, 2020.

596

597 Elissa Mhanna and Mohamad Assaad. Countering the communication bottleneck in federated
 598 learning: A highly efficient zero-order optimization technique. *Journal of Machine Learning
 599 Research*, 25(418):1–53, 2024.

600 Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. The e2e dataset: New challenges for
 601 end-to-end generation. *arXiv preprint arXiv:1706.09254*, 2017.

602

603 Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012
 604 IEEE conference on computer vision and pattern recognition*, pp. 3498–3505. IEEE, 2012.

605

606 A Paszke. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint
 607 arXiv:1912.01703*, 2019.

608

609 Bowen Ping, Shuo Wang, Hanqing Wang, Xu Han, Yuzhuang Xu, Yukun Yan, Yun Chen, Baobao
 610 Chang, Zhiyuan Liu, and Maosong Sun. Delta-come: Training-free delta-compression with
 611 mixed-precision for large language models. *arXiv preprint arXiv:2406.08903*, 2024.

612

613 Zeju Qiu, Weiyang Liu, Haiwen Feng, Yuxuan Xue, Yao Feng, Zhen Liu, Dan Zhang, Adrian Weller,
 614 and Bernhard Schölkopf. Controlling text-to-image diffusion by orthogonal finetuning. *Advances
 615 in Neural Information Processing Systems*, 36:79320–79362, 2023.

616

617 Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language
 618 understanding by generative pre-training. 2018.

619

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

622

623 Mohaimenul Azam Khan Raiaan, Md Saddam Hossain Mukta, Kaniz Fatema, Nur Mohammad
 624 Fahad, Sadman Sakib, Most Marufatul Jannat Mim, Jubaer Ahmad, Mohammed Eunus Ali, and
 625 Sami Azam. A review on large language models: Architectures, applications, taxonomies, open
 626 issues and challenges. *IEEE access*, 12:26839–26874, 2024.

627

628 Google Research. Vision Transformer Models on Hugging Face. <https://huggingface.co/google>, 2025. Accessed: 06-May-2025.

629

630 Maria Seleznova, Dana Weitzner, Raja Giryes, Gitta Kutyniok, and Hung-Hsu Chou. Neural (tangent
 631 kernel) collapse. *Advances in Neural Information Processing Systems*, 36:16240–16270, 2023.

632

633 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy
 634 Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.

635

636 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
 637 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
 638 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

639

640 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 641 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing
 642 systems*, 30, 2017.

643

644 Shilong Wang, Jianchun Liu, Hongli Xu, Jiaming Yan, and Xianjun Gao. Efficient federated fine-
 645 tuning of large language models with layer dropout. *arXiv preprint arXiv:2503.10217*, 2025.

646

647 Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and
 Pietro Perona. Caltech-ucsd birds 200. 2010.

648

649 Ross Wightman. timm: PyTorch Image Models. <https://github.com/huggingface/pytorch-image-models/tree/main/timm>, 2025. Accessed: 06-May-2025.

650

651 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
 652 Pierrick Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art
 653 natural language processing. In *Proceedings of the 2020 conference on empirical methods in
 654 natural language processing: system demonstrations*, pp. 38–45, 2020.

648 Taiqiang Wu, Jiahao Wang, Zhe Zhao, and Ngai Wong. Mixture-of-subspaces in low-rank adaptation.
 649 *arXiv preprint arXiv:2406.11909*, 2024.
 650

651 Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database:
 652 Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on*
 653 *computer vision and pattern recognition*, pp. 3485–3492. IEEE, 2010.

654 Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. Parameter-efficient
 655 fine-tuning methods for pretrained language models: A critical review and assessment. *arXiv*
 656 *preprint arXiv:2312.12148*, 2023a.
 657

658 Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen,
 659 Xiaopeng Zhang, and Qi Tian. Qa-lora: Quantization-aware low-rank adaptation of large language
 660 models. *arXiv preprint arXiv:2309.14717*, 2023b.
 661

662 Prateek Yadav, Leshem Choshen, Colin Raffel, and Mohit Bansal. Compeft: Compression for
 663 communicating parameter efficient updates via sparsification and quantization. *arXiv preprint*
 664 *arXiv:2311.13171*, 2023.
 665

666 Xiaozhe Yao, Qinghao Hu, and Ana Klimovic. Deltazip: Efficient serving of multiple full-model-
 667 tuned llms. In *Proceedings of the Twentieth European Conference on Computer Systems*, pp.
 668 110–127, 2025.
 669

670 Chunxu Zhang, Guodong Long, Tianyi Zhou, Zijian Zhang, Peng Yan, and Bo Yang. When federated
 671 recommendation meets cold-start problem: Separating item attributes and user interactions. In
 672 *Proceedings of the ACM Web Conference 2024*, pp. 3632–3642, 2024.
 673

674 Yuanhe Zhang, Fanghui Liu, and Yudong Chen. Lora-one: One-step full gradient could suffice for
 675 fine-tuning large language models, provably and efficiently. *arXiv preprint arXiv:2502.01235*,
 676 2025a.
 677

678 Yuanhe Zhang, Fanghui Liu, and Yudong Chen. One-step full gradient suffices for low-rank fine-
 679 tuning, provably and efficiently. *arXiv preprint arXiv:2502.01235*, 2025b.
 680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

APPENDIX

A PROOF OF THEOREM 1

Let $\Delta W^* \in \mathbb{R}^{m \times n}$ denote the optimal adapter for the downstream task, ΔW the adapter obtained by LoRA fine-tuning, and $\tilde{\Delta W}$ the SOLAR reconstruction. Let ΔW_{proj} denote the projection of ΔW onto the SOLAR bases (i.e., bases that are constructed from the SVD of the foundation model's weights, combined with randomized perturbations).

Our proof relies on the following standard assumptions from the literature on parameter-efficient fine-tuning and randomized numerical linear algebra:

- (A1) *Spectral Initialization*: The LoRA adapter matrices A and B are initialized using the spectral initialization strategy from Zhang et al. (2025a).
- (A2) *Low-Rank Update*: The optimal task-specific update ΔW^* is approximately low-rank, with rank $r^* < \min\{m, n\}$ Zhang et al. (2025a).
- (A3) *Well-Behaved Data*: The training data follows the generation process outlined in Zhang et al. (2025a), where input features are drawn from an isotropic sub-Gaussian or Gaussian distribution.
- (A4) *Fast Spectrum Decay*: The projected update matrix ΔW_{proj} exhibits spectral decay, meaning its tail singular values are small (Martinsson & Tropp, 2020).

First, we decompose the total error using the triangle inequality. The total error, $\|\Delta \tilde{W} - \Delta W^*\|_F$, is the distance between the SOLAR-reconstructed adapter and the optimal adapter. This is bounded by the sum of the Training Error and the Compression Error:

$$\|\Delta \tilde{W} - \Delta W^*\|_F \leq \underbrace{\|\Delta \tilde{W} - \Delta W\|_F}_{\text{Compression Error}} + \underbrace{\|\Delta W - \Delta W^*\|_F}_{\text{Training Error}} \quad (6)$$

Here, the first term, $\|\Delta \tilde{W} - \Delta W\|_F$, is the compression error introduced by SOLAR's approximation. The second term, $\|\Delta W - \Delta W^*\|_F$, is the training error from the underlying LoRA fine-tuning process itself. We will bound each term separately.

The analysis of the training error for LoRA adapters is non-trivial and has been extensively studied. We directly leverage the results from Zhang et al. (2025a), showing that under Assumptions (A1)-(A3), LoRA trained with gradient descent converges to the optimal low-rank adapter ΔW^* . Their analysis provides the following bound on the training error after t steps:

$$\|\Delta W - \Delta W^*\|_F \leq \sqrt{2r^*} \left(1 - \frac{\eta \lambda_{r^*}}{64\kappa}\right)^t \lambda_{r^*}, \quad (7)$$

where r^* is the rank of the optimal update ΔW^* , κ is its condition number, λ_{r^*} is its r^* -th singular value, and η is the learning rate. This bound, derived under the specified spectral initialization and data concentration assumptions, demonstrates that the fine-tuned adapter ΔW gets exponentially closer to the optimal adapter ΔW^* as training progresses.

SOLAR reconstructs the adapter as a sparse coefficientization over these perturbed bases:

$$\Delta \tilde{W} = \sum_{i=1}^{N_B} \sum_{j=1}^{N_A} \beta_i \alpha_j M_B^{(i)} M_A^{(j)}. \quad (8)$$

Following the randomized rangefinder formulation Halko et al. (2011); Martinsson & Tropp (2020), we construct the sketch matrices for both the column and row spaces of the LoRA-style adapter update ΔW as

$$Y_A = \Delta W \Omega_A \in \mathbb{R}^{m \times N_A}, \quad Y_B = \Delta W^\top \Omega_B \in \mathbb{R}^{n \times N_B}. \quad (9)$$

Each column of Y_A represents the action of ΔW on a random probe vector drawn from the right-basis pool Ω_A , effectively sampling the column space of ΔW . Similarly, each column of Y_B captures

random projections of the row space of ΔW . These sketches compactly encode the dominant directions of ΔW without explicitly computing its singular value decomposition.

The Gaussian perturbations in $M_A^{(i)} = V_{:, \mathcal{I}_i} + \epsilon_i$ and $M_B^{(j)} = U_{:, \mathcal{J}_j} + \epsilon_j$ play an important theoretical and practical role. First, they ensure that the composite sketching matrices Ω_A and Ω_B satisfy the sub-Gaussian concentration and Johnson–Lindenstrauss properties required for the probabilistic error bounds in randomized numerical linear algebra Halko et al. (2011). Second, adding small isotropic noise expands the effective span of the sampled singular directions, preventing over-alignment with any single dominant mode and improving numerical stability when the singular spectrum of ΔW decays slowly. Finally, this perturbation acts as a regularizer that mitigates sampling bias inherited from the foundation model’s specific singular subspace, ensuring broader coverage of the subspace where fine-tuned updates lie.

We then compute orthonormal bases for the column spans of these sketches:

$$Q_A = \text{orth}(Y_A) \in \mathbb{R}^{m \times q_A}, \quad Q_B = \text{orth}(Y_B) \in \mathbb{R}^{n \times q_B}, \quad (10)$$

where

$$r_A = \text{rank}(Q_A) \leq \min(m, N_A), \quad r_B = \text{rank}(Q_B) \leq \min(n, N_B).$$

By construction, $\text{range}(Q_A) = \text{range}(Y_A)$ and $\text{range}(Q_B) = \text{range}(Y_B)$. In the terminology of randomized numerical linear algebra, this process corresponds to the *rangefinder step*, which identifies low-dimensional subspaces that approximate the dominant column and row spaces of ΔW .

Finally, we define the two-sided (bi-rangefinder) projection as

$$\mathcal{P}_{N_A, N_B}(\Delta W) := Q_A Q_A^\top \Delta W Q_B Q_B^\top. \quad (11)$$

This projection provides a low-rank approximation to ΔW using orthonormal subspaces inferred from randomized sketches. Geometrically, $\mathcal{P}_{N_A, N_B}(\Delta W)$ captures the principal subspace of ΔW identified by Ω_A and Ω_B , offering an efficient surrogate for the optimal SVD-based projection $U_1 U_1^\top \Delta W V_1 V_1^\top$ while retaining probabilistic error guarantees Halko et al. (2011); Martinsson & Tropp (2020).

We bound the bi-projection error by splitting it into two one-sided parts using projector non-expansiveness ($\|Q_A Q_A^\top X\|_F \leq \|X\|_F$):

$$\begin{aligned} \|\Delta W - Q_A Q_A^\top \Delta W Q_B Q_B^\top\|_F &\leq \|\Delta W - Q_A Q_A^\top \Delta W\|_F + \|Q_A Q_A^\top (\Delta W - \Delta W Q_B Q_B^\top)\|_F \\ &\leq \|\Delta W - Q_A Q_A^\top \Delta W\|_F + \|\Delta W - \Delta W Q_B Q_B^\top\|_F. \end{aligned} \quad (12)$$

Each addend is a standard one-sided rangefinder error. By Theorem 10.5 of Halko et al. (2011) (Frobenius form) with oversampling $N_A > r_A + 1$ and $N_B > r_B + 1$,

$$\mathbb{E} \|\Delta W - Q_A Q_A^\top \Delta W\|_F \leq \left(1 + \frac{r_A}{N_A - r_A - 1}\right)^{\frac{1}{2}} \left(\sum_{t > r_A} \sigma_t(\Delta W)^2\right)^{\frac{1}{2}}, \quad (13)$$

$$\mathbb{E} \|\Delta W - \Delta W Q_B Q_B^\top\|_F \leq \left(1 + \frac{r_B}{N_B - r_B - 1}\right)^{\frac{1}{2}} \left(\sum_{t > r_B} \sigma_t(\Delta W)^2\right)^{\frac{1}{2}}. \quad (14)$$

Combining equation 12–equation 14 yields the expected two-sided projection error bound:

$$\mathbb{E} \|\Delta W - \mathcal{P}_{N_A, N_B}(\Delta W)\|_F \leq \left(1 + \frac{r_A}{N_A - r_A - 1}\right)^{\frac{1}{2}} \left(\sum_{t > r_A} \sigma_t^2\right)^{\frac{1}{2}} + \left(1 + \frac{r_B}{N_B - r_B - 1}\right)^{\frac{1}{2}} \left(\sum_{t > r_B} \sigma_t^2\right)^{\frac{1}{2}}. \quad (15)$$

(When desired, power iterations can be incorporated on either side to sharpen the spectral decay and constants Halko et al. (2011); Martinsson & Tropp (2020).)

After projection, SOLAR enforces sparsity by retaining only the top- k basis pairs in equation 8. Let the singular values of $\mathcal{P}_{N_A, N_B}(\Delta W)$ be $\{\tilde{\sigma}_t\}$, we have:

$$\|\Delta \tilde{W} - \mathcal{P}_{N_A, N_B}(\Delta W)\|_F \leq \left(\sum_{t > k} \tilde{\sigma}_t^2\right)^{\frac{1}{2}}. \quad (16)$$

Moreover, orthogonal projections are contractions in Frobenius norm and cannot increase tail energy, hence

$$\sum_{t>k} \tilde{\sigma}_t^2 \leq \sum_{t>k} \sigma_t(\Delta W)^2. \quad (17)$$

Adding and subtracting $\mathcal{P}_{N_A, N_B}(\Delta W)$ and using equation 15–equation 17, we obtain

$$\begin{aligned} \mathbb{E} \|\Delta \tilde{W} - \Delta W\|_F &\leq \mathbb{E} \|\Delta W - \mathcal{P}_{N_A, N_B}(\Delta W)\|_F + \mathbb{E} \|\Delta \tilde{W} - \mathcal{P}_{N_A, N_B}(\Delta W)\|_F \\ &\leq \left(1 + \frac{r_A}{N_A - r_A - 1}\right)^{\frac{1}{2}} \left(\sum_{t>r_A} \sigma_t^2\right)^{\frac{1}{2}} + \left(1 + \frac{r_B}{N_B - r_B - 1}\right)^{\frac{1}{2}} \left(\sum_{t>r_B} \sigma_t^2\right)^{\frac{1}{2}} \end{aligned} \quad (18)$$

$$+ \left(\sum_{t>k} \sigma_t^2\right)^{\frac{1}{2}}. \quad (19)$$

Combining the decomposition with equation 19 and the LoRA training bound equation 7, we conclude

$$\begin{aligned} \mathbb{E} \|\Delta \tilde{W} - \Delta W^*\|_F &\leq \underbrace{\left(1 + \frac{r_A}{N_A - r_A - 1}\right)^{\frac{1}{2}} \left(\sum_{t>r_A} \sigma_t^2\right)^{\frac{1}{2}} + \left(1 + \frac{r_B}{N_B - r_B - 1}\right)^{\frac{1}{2}} \left(\sum_{t>r_B} \sigma_t^2\right)^{\frac{1}{2}}}_{\text{projection error}} \\ &\quad + \underbrace{\left(\sum_{t>k} \sigma_t^2\right)^{\frac{1}{2}}}_{\text{sparsification error}} + \underbrace{\sqrt{2r^*} \left(1 - \frac{\eta \lambda_{r^*}}{64\kappa}\right)^t \lambda_{r^*}}_{\text{training error}}. \end{aligned} \quad (20)$$

Each term in equation 20 can be driven to zero under mild conditions: (i) the projection error vanishes as N_A, N_B grow so that r_A, r_B reach the true (or effective) rank of ΔW (then the corresponding spectral tails are zero); (ii) the sparsification error vanishes when k exceeds the numerical rank of $\mathcal{P}_{N_A, N_B}(\Delta W)$; and (iii) the training error decays to zero as $t \rightarrow \infty$ under (A1)–(A3) by equation 7. Consequently, with sufficient sampling (N_A, N_B), sparsity budget (k), $\mathbb{E} \|\Delta \tilde{W} - \Delta W^*\|_F \rightarrow 0$.

B IMPLEMENTATION DETAILS

All models are implemented using PyTorch Paszke (2019), with HuggingFace Transformers Wolf et al. (2020) for LLaMA and GPT-based models, and Timm Wightman (2025) for ViT-based vision backbones. Training and evaluation are performed on NVIDIA A100 and RTX 4090 GPUs. For all vision experiments, we use ViT-B and ViT-L as base encoders. For language models, we use GPT-2 and LLaMA-3 (1B, 3B, 8B). LoRA is applied to the query and value projections. SOLAR operates post-training by compressing the PEFT adapter matrices. All experiments are conducted under a fixed random seed for reproducibility. The implementation code for SOLAR, along with scripts used to reproduce the experiments, is included in the supplementary material and also available at <https://anonymous.4open.science/r/SOLAR-D3B2/>.

C DATASET DETAILS

We summarize dataset statistics in Table 8, including number of training samples and class counts.

We summarize dataset statistics used in the LLM experiments in Table 9, covering instruction tuning (Section 3.2) and language generation tasks (Section 3.3). The table includes the number of training samples, average sequence lengths, and the model-specific context in which each dataset is used in the experiments.

864 Table 8: Dataset statistics used in experiments. Each dataset includes the number of training samples
 865 and classes.

867	Dataset	Training Samples	Number of Classes
868	CIFAR-10	50,000	10
869	CIFAR-100	50,000	100
870	Food-101	75,750	101
871	Tiny-ImageNet	100,000	200
872	ImageNet-1K	1,281,167	1,000

874
875 Table 9: Dataset statistics in LLM experiments.

877	Dataset	Samples	Avg. Seq. Length	Context
878	Stanford Alpaca	52,000	~256 tokens	LLaMA-3 instruction tuning
879	MMLU	15,858	~200 tokens	LLaMA-3 Generalization evaluation
880	E2E NLG	42,000	~35 tokens	GPT-2 generation fine-tuning

883 D REPRESENTATION COST DETAILS: PARAMETERS AND STORAGE

884
 885 To quantify SOLAR’s compression benefit, we detail the number of adapter parameters and byte-level
 886 footprint across ViT-B, ViT-L, LLaMA, and GPT-2 models. We compare LoRA, NOLA, and SOLAR
 887 under adapter rank ($r = 4$). Tables 10 through 15 provide full parameter breakdowns. Byte-level
 888 analysis is presented in Table 13.

889
 890 **ViT.** For vision backbones, Table 10 and Table 11 report the number of representation parameters
 891 for query projections (Q) and classifier heads. In the experiments presented in the main paper, the
 892 classifier head parameters are excluded from comparison since they are identical across all methods
 893 following Koohpayegani et al. (2024). NOLA’s parameter footprint for MLP projections is shown in
 894 Table 12 (following the setup in Koohpayegani et al. (2024)). Byte-level storage comparisons across
 895 quantization, used to produce Table 2 and Table 3 in the main paper, are provided in Table 13.

896
 897 Table 10: Number of representation parameters for ViT-B (Rank = 4). Each row reports the parameter
 898 count for query projections and the classifier head using SOLAR and LoRA across different datasets.
 899 The classifier head parameter count is shared across methods and is computed as $(\text{num_classes} \times 768 + \text{num_classes})$. For SOLAR, the query projection count corresponds to: number of
 900 layers \times (top _{k} coefficients for A + top _{k} coefficients for B + encoded basis for A + encoded basis for
 901 B) + 1 (seed value). All SOLAR rows follow the form $N \rightarrow \text{top}_k$ where N is the original subspace
 902 size. For LoRA, the query projection count corresponds to: number of layers \times (input dimension \times
 903 rank for A + rank \times output dimension for B), where rank is 4.

906	Method	Dataset	Query (Q)	Classifier Head
907 908 909 910	SOLAR	CIFAR-10	$12 \times ((1600 + 1600) + \frac{4000+4000}{32}) + 1 = 41,401$	$10 \times 768 + 10 = 7,690$
		CIFAR-100	41,401	$100 \times 768 + 100 = 76,900$
		Food-101	41,401	$101 \times 768 + 101 = 77,669$
		Tiny-ImageNet	41,401	$200 \times 768 + 200 = 154,000$
911 912 913	LoRA	CIFAR-10	$12 \times [(768 \times 4) + (4 \times 768)] = 73,728$	$10 \times 768 + 10 = 7,690$
		CIFAR-100	73,728	$100 \times 768 + 100 = 76,900$
		Food-101	73,728	$101 \times 768 + 101 = 77,669$
		Tiny-ImageNet	73,728	$200 \times 768 + 200 = 154,000$

914
 915 **LLMs.** For language models, parameter counts for adapter layers are detailed in Table 14 for
 916 LLaMA and in Table 15 for GPT-2 variants.

918
 919 Table 11: Number of representation parameters for ViT-L (Rank = 4). Each row shows the parameter
 920 counts for Query projections and the classifier head using SOLAR and LoRA across different datasets.
 921 The classifier head parameter count is shared across methods and is calculated as (num_classes
 922 $\times 1024 + \text{num_classes}$).
 923

Method	Dataset	Query (Q)	Classifier Head
SOLAR	CIFAR-10	$24 \times ((500 + 500) + \frac{1000+1000}{32}) + 1 = 25,501$	$10 \times 1024 + 10 = 10,250$
	CIFAR-100	25,501	$100 \times 1024 + 100 = 102,500$
	Food-101	25,501	$101 \times 1024 + 101 = 103,625$
	Tiny-ImageNet	25,501	$200 \times 1024 + 200 = 204,800$
LoRA	CIFAR-10	$24 \times [(1024 \times 4) + (4 \times 1024)] = 196,608$	$10 \times 1024 + 10 = 10,250$
	CIFAR-100	196,608	$100 \times 1024 + 100 = 102,500$
	Food-101	196,608	$101 \times 1024 + 101 = 103,625$
	Tiny-ImageNet	196,608	$200 \times 1024 + 200 = 204,800$

924
 925
 926
 927
 928
 929
 930
 931
 932 Table 12: Number of representation parameters for ViT-B (Rank = 4). Each row shows the pa-
 933 rameter counts for MLP projections (for NOLA) and classifier head across datasets. The classifier
 934 head parameter count is shared across methods and is calculated as (num_classes $\times 768 +$
 935 num_classes).
 936

Method	Dataset	MLP	Classifier Head
NOLA	CIFAR-10	$12 \times 2 \times 2 \times 1000 + 1 = 48,001$	$10 \times 768 + 10 = 7,690$
	CIFAR-100	48,001	$100 \times 768 + 100 = 76,900$
	Food-101	48,001	$101 \times 768 + 101 = 77,669$
	Tiny-ImageNet	48,001	$200 \times 768 + 200 = 154,000$

E ADDITIONAL EXPERIMENTAL RESULTS

943
 944 This section provides supplementary experimental results to further validate the claims made in the
 945 main paper. We present detailed performance metrics for additional model scales and include a
 946 crucial ablation study that compares SOLAR against a parameter-matched LoRA baseline.
 947

E.1 PERFORMANCE ON INTERMEDIATE-SCALE LLaMA MODELS

948
 949 Table 16 extends our analysis to the LLaMA-3.2 3B and LLaMA-3.1 8B models, demonstrating
 950 SOLAR’s consistent efficiency and performance on intermediate-scale architectures. The results show
 951 that SOLAR maintains the performance of the original LoRA adapters while achieving parameter
 952 reductions of over 90%.
 953

E.2 COMPRESSION OF ADAPTIVE-RANK PEFT METHODS (ADALORA)

954
 955 To evaluate SOLAR on more recent PEFT methods, we applied it to AdaLoRA, which produces
 956 adaptive-rank adapter matrices (\mathbf{A} and \mathbf{B}). SOLAR compresses these trained adapters post-hoc, using
 957 an initial rank of $r = 8$ and a target average rank of $r = 1$ on LLaMA-3.2 3B and LLaMA-2 13B.
 958 As shown in Table 17, SOLAR significantly reduces adapter parameters while preserving MMLU
 959 performance.
 960

E.2.1 EXPERIMENTS WITH 2-BIT QUANTIZATION

961
 962 To further validate SOLAR’s robustness to aggressive quantization, we conducted additional experi-
 963 ments with 2-bit quantization on LLaMA-2 13B and LLaMA-3.1 8B. The results, summarized in
 964 Table 18, confirm that SOLAR remains effective while drastically reducing parameter counts.
 965

E.3 EXTREME COMPRESSION

966
 967 In this section, we report additional experiments demonstrating SOLAR’s ability to achieve extreme
 968 compression while retaining competitive accuracy. These results complement the main paper by
 969

972
 973 Table 13: Byte-level footprint of representation parameters for ViT-B and ViT-L using LoRA and
 974 SOLAR. Each value reflects the total number of bytes required to store adapter updates (excluding
 975 classifier heads). For LoRA, storage is computed as: number of layers \times (rank \times output dimension
 976 for B + input dimension \times rank for A) \times precision in bytes (e.g., 4 bytes for 32-bit float). For
 977 SOLAR, storage is computed as: number of layers \times (top $_k$ coefficients for A + top $_k$ coefficients
 978 for B + encoded basis vectors for A + encoded basis for B) \times precision in bytes, plus 1 byte to
 979 store a random seed. For example, the row "500 \rightarrow 50" denotes that 500-dimensional subspaces are
 980 sparsified to top- k = 50 coefficients, with encoded bases represented at 1 bit per element (8 elements
 981 per byte).

Method	Representation Footprint (Bytes)
LoRA ($r=1$)	$12 \times [(768 \times 1) + (1 \times 768)] \times 4 = 73,728$
SOLAR for ViT-B 8Bit ($r=1$, 500 \rightarrow 50)	$12 \times [(50 + 50) + \frac{500}{8}] \times 1 + 1 = 1,951$
SOLAR for ViT-B 8Bit ($r=1$, 100 \rightarrow 10)	$12 \times [(10 + 10) + \frac{100}{8}] \times 1 + 1 = 391$
LoRA ($r=4$)	$24 \times [(1024 \times 4) + (4 \times 1024)] \times 4 = 786,432$
SOLAR for ViT-L 32Bit ($r=4$, 4000 \rightarrow 1600)	$24 \times [(1600 + 1600) + \frac{4000}{32}] \times 4 + 1 = 319,201$
SOLAR for ViT-L 16Bit ($r=4$, 4000 \rightarrow 1600)	$24 \times [(1600 + 1600) + \frac{4000}{16}] \times 2 + 1 = 165,601$
SOLAR for ViT-L 8Bit ($r=4$, 4000 \rightarrow 1600)	$24 \times [(1600 + 1600) + \frac{4000}{8}] \times 1 + 1 = 88,801$
SOLAR for ViT-L 4Bit ($r=4$, 4000 \rightarrow 1600)	$24 \times [(1600 + 1600) + \frac{4000}{4}] \times 0.5 + 1 = 50,401$

992
 993
 994
 995 Table 14: Number of representation parameters for LLaMA-3 models using LoRA, NOLA, and
 996 SOLAR. Each row reports total adapter parameters for attention projections (Q and V for LoRA
 997 and NOLA; Q and K for SOLAR). Output heads and MLP layers are frozen. For LoRA, the
 998 parameter count is computed as: number of layers \times (input dimension \times rank for B + rank \times output
 999 dimension for A +). Due to differing dimensions between A and B in LoRA, the table computes the
 1000 contributions for Q and V projections separately. For NOLA, it is computed as: number of layers
 1001 \times 2 \times (number of random basis vectors), assuming separate basis sets for A and B . For SOLAR,
 1002 the count is: number of layers \times 2 \times (top $_k$ coefficients for B + top $_k$ for A + encoded bases for B +
 1003 encoded bases for A), plus 1 byte to communicate or store the shared seed.

Model (Rank)	Configuration	Total Parameters
LLaMA-3.2 1B ($r=8$)	16 layers (Q, V)	$16 \times [(2048 \times 8 + 8 \times 2048) + (2048 \times 8 + 8 \times 512)] = 851,968$
NOLA	16 layers (Q, V)	$16 \times 2 \times (1000 + 1000) = 64,000$
SOLAR ($r=8, 4K \rightarrow 1.2K$)	16 layers (Q, V)	$16 \times 2 \times (1200 + 1200 + \frac{4000}{32}) + 1 = 80,801$
LLaMA-3.2 3B ($r=1$)	28 layers (Q, V)	$28 \times [(3072 \times 1 + 1 \times 3072) + (3072 \times 1 + 1 \times 1024)] = 286,720$
NOLA	28 layers (Q, V)	$28 \times 2 \times (1000 + 1000) = 112,000$
SOLAR ($r=1, 1000 \rightarrow 150$)	28 layers (Q, V)	$28 \times 2 \times (150 + 150 + \frac{1000}{32}) + 1 = 18,551$
LLaMA-3.1 8B ($r=1$)	32 layers (Q, V)	$32 \times [(4096 \times 1 + 1 \times 4096) + (4096 \times 1 + 1 \times 1024)] = 425,984$
NOLA	32 layers (Q, V)	$32 \times 2 \times (1000 + 1000) = 128,000$
SOLAR ($r=1, 1000 \rightarrow 300$)	32 layers (Q, V)	$32 \times 2 \times (300 + 300 + \frac{1000}{32}) + 1 = 40,401$

1014
 1015
 1016
 1017
 1018
 1019 highlighting scenarios where communication and storage constraints are especially strict (e.g.,
 1020 distributed or on-device learning).

1021 Table 19 shows evaluations on four vision datasets using ViT-B under different compression budgets.
 1022 We quantify the bit-level representation footprint assuming 32-bit precision during training and apply
 1023 8-bit quantization to the SOLAR coefficients after top- k selection. Compared to LoRA ($r = 1$),
 1024 SOLAR reduces the adapter footprint by up to 99% (from 74KB to 0.4KB) with only minor drops in
 1025 accuracy. These results illustrate that SOLAR enables fine-grained tradeoffs between accuracy and
 storage cost under extreme compression budgets.

1026
 1027 Table 15: Number of trainable adapter parameters for GPT-2 models using LoRA, NOLA, and
 1028 SOLAR. Each row reports the total number of parameters added to the query and value projections (Q
 1029 and V). All configurations freeze the output heads and MLP layers. For LoRA, the parameter count is
 1030 computed as: number of layers $\times 2 \times (\text{input dimension} \times \text{rank for } B + \text{rank} \times \text{output dimension for } A)$. For NOLA, the parameter count is: number of layers $\times 2 \times (\text{number of random basis vectors})$,
 1031 assuming separate basis sets for Q and V. For SOLAR, the parameter count is: number of layers $\times 2 \times (\text{top}_k \text{ coefficients for } B + \text{top}_k \text{ coefficients for } A + \text{encoded bases for } B + \text{encoded bases for } A)$,
 1032 plus 1 for the shared seed.
 1033

Model (Rank)	Configuration	Total Parameters
GPT-2 Small ($r=4$)	12 layers (Q, V)	$12 \times 2 \times (768 \times 4 + 4 \times 768) = 147,456$
NOLA	12 layers (Q, V)	$12 \times 2 \times (1000 + 1000) = 48,000$
SOLAR ($r=1, 1000 \rightarrow 300$)	12 layers (Q, V)	$12 \times 2 \times (300 + 300 + \frac{1000}{32}) + 1 = 15,150$
SOLAR ($r=1, 100 \rightarrow 90$)	12 layers (Q, V)	$12 \times 2 \times (90 + 90 + \frac{100}{32}) + 1 = 4,396$
GPT-2 Medium ($r=4$)	24 layers (Q, V)	$24 \times 2 \times (1024 \times 4 + 4 \times 1024) = 393,216$
NOLA	24 layers (Q, V)	350,000 Koohpayegani et al. (2024)
SOLAR ($r=4, 1000 \rightarrow 300$)	24 layers (Q, V)	$24 \times 2 \times (300 + 300 + \frac{1000}{32}) + 1 = 30,301$
SOLAR ($r=4, 100 \rightarrow 90$)	24 layers (Q, V)	$24 \times 2 \times (90 + 90 + \frac{100}{32}) + 1 = 8,791$

1045
 1046 Table 16: Model representation efficiency for LLaMA 3B and 8B models. For the 8B model, all
 1047 methods use 4-bit quantization, making the LoRA baseline equivalent to QLoRA.
 1048

Model	LLaMA-3.2 3B			LLaMA-3.1 8B (4-bit)		
	LoRA $r=1$	NOLA 1000 bases	SOLAR $r=1(1K \rightarrow 0.1K)$	LoRA $r=1$	NOLA 1000 bases	SOLAR $r=1(1K \rightarrow 0.3K)$
# Params	287K	112K	16K (94% ↓)	425K	128K	40K (91% ↓)
Val Loss	1.02	1.31	1.04	0.89	1.01	0.90
MMLU Acc	54.0	52.7	54.0	60.9	56.1	60.9

F SCALABILITY TO LARGER VISION MODELS

1059 To validate that SOLAR remains effective and computationally tractable on larger-scale models, we
 1060 conducted experiments on the ViT-G/14 architecture. This model is substantially larger than the
 1061 ViT-B/L backbones used in our main experiments, providing a strong test of scalability.
 1062

1063 We fine-tuned a ViT-G/14 model on the full CIFAR-10, CIFAR-100, Food-101, and T-ImageNet
 1064 datasets using a LoRA adapter with rank $r = 4$. We then applied SOLAR with a basis pool of 8,000
 1065 vectors, selecting the top 4,000 coefficients to form the compressed adapter.
 1066

1067 As shown in Table 20, SOLAR successfully preserves the performance of the original LoRA adapter
 1068 with negligible accuracy drops, while reducing the adapter’s parameter count by 31% (from 492K to
 1069 340K). This result demonstrates that SOLAR’s core mechanisms—including SVD extraction and
 1070 sparse reconstruction—scale effectively to larger models without sacrificing compression efficiency
 or task performance.
 1071

F.1 ABLATION STUDY: BUDGET-MATCHED LORA COMPARISON

1073 To further validate the efficiency of our compression strategy, we conduct an ablation study directly
 1074 comparing SOLAR to a budget-matched LoRA baseline, as suggested by reviewer feedback.[1] This
 1075 comparison is critical to demonstrate that SOLAR’s benefits extend beyond mere parameter reduction
 1076 and offer a more effective performance-compression trade-off than simply training a lower-rank
 1077 adapter from scratch.
 1078

1079 As shown in Table 21, fine-tuning a LoRA adapter with a reduced rank ($r=2$) to match the parameter
 1080 count of the compressed SOLAR adapter results in a significant performance degradation across all
 1081

1080 Table 17: SOLAR applied to AdaLoRA adapters on intermediate-scale LLaMA models.
1081

Method	# Params (Adapter)	MMLU Accuracy
AdaLoRA (Baseline, 3B)	305K	54.8%
SOLAR (on AdaLoRA, 3B)	16K	54.7%
AdaLoRA (Baseline, 13B)	871K	57.9%
SOLAR (on AdaLoRA, 13B)	16K	57.7%

1088 Table 18: 2-bit quantization experiments comparing LoRA (QLoRA) and SOLAR.
1089

Method	Quantization	# Params	MMLU Acc
LoRA (QLoRA) - LLaMA-2 13B	2-bit	410K	53.1
SOLAR _{r=1(1K→0.3K)} - LLaMA-2 13B	2-bit	51K	53.1
LoRA (QLoRA) - LLaMA-3.1 8B	2-bit	363K	58.4
SOLAR _{r=1(1K→0.3K)} - LLaMA-3.1 8B	2-bit	40K	58.4

1097 tasks. In contrast, SOLAR, when applied to the higher-performing LoRA ($r=4$) adapter, successfully
1098 preserves task accuracy while achieving a comparable parameter budget. This highlights that SOLAR
1099 retains the expressive power of the original higher-rank adapter, a feat not achievable by simply
1100 reducing the rank during training. All experiments were conducted on the full datasets using the
1101 ViT-B backbone, with results reported as the mean accuracy over five independent runs to ensure
1102 statistical robustness.

1104 G COMPARISON WITH SIMPLE SVD TRUNCATION

1107 To compare against simple post-hoc SVD truncation, we evaluate SOLAR’s performance against
1108 SVD applied directly to the LoRA update ΔW . Since the LoRA adapter ΔW already has rank r ,
1109 SVD only provides compression if the truncation rank is set lower than r . We use an initial LoRA
1110 rank of $r = 4$ and truncate the SVD to rank 1. In contrast, SOLAR achieves a much smaller footprint
1111 by reparameterizing the update in the foundation model’s subspace. The results are summarized in
1112 Table 22.

1113 H APPLICATION TO FEDERATED LEARNING

1116 One of the motivations for developing SOLAR is to reduce communication overhead in distributed
1117 learning scenarios, such as Federated Learning (FL). In typical FL setups, clients fine-tune a model
1118 on their local data and transmit the resulting model updates (e.g., LoRA adapters) to a central server
1119 for aggregation. As highlighted by recent work Mhanna & Assaad (2024), communication—not
1120 computation—is often the primary bottleneck. Transmitting full adapters from thousands of clients
1121 can generate enormous data transfer loads. For example, in an FL setup with 10,000 clients—1,000
1122 participating in each of 10 training rounds—transmitting 74 KB LoRA adapters per client would
1123 amount to 740 GB of total data transfer.

1124 SOLAR addresses this challenge as a lightweight, post-hoc compression utility. After local training,
1125 each client can compress its adapter with SOLAR before transmission. The server then receives only
1126 the sparse coefficients and a random seed, drastically reducing per-client communication costs.

1127 To demonstrate SOLAR’s effectiveness in distributed settings, we simulated a 10-client FL environment.
1128 We compare a baseline where clients transmit full LoRA adapters with a scenario where clients
1129 transmit SOLAR-compressed adapters. Each client fine-tunes a ViT-B model on CIFAR-10 with
1130 LoRA ($r = 4$), under two data distribution scenarios: an IID baseline and a non-IID distribution
1131 generated via a Dirichlet process with a concentration parameter of 0.5. The simulation runs for 30
1132 communication rounds, with one epoch of local training per client per round.

1133 As shown in Table 23, the performance gap between full LoRA adapters and SOLAR-compressed
1134 adapters is minimal in both IID and non-IID settings. This demonstrates that SOLAR’s compression

1134 Table 19: Evaluation of extreme compression on ViT-B. We report bit-level representation footprint
 1135 (32-bit baseline) and top-1 accuracy over 5 runs. All models are trained for 10 epochs.
 1136

Method	Byte Footprint	Oxford Pets	SUN397	CUB-200	ImageNet-1K
LoRA ($r=1$)	74KB	93.0 ± 0.5	74.3 ± 0.3	84.7 ± 0.4	81.5 ± 0.6
SOLAR ($r=1, 500 \rightarrow 50$)	<u>2KB</u> (97% \downarrow)	91.2 ± 0.6	72.4 ± 0.4	81.4 ± 0.5	80.7 ± 0.4
SOLAR ($r=1, 100 \rightarrow 10$)	0.4KB (99% \downarrow)	90.3 ± 0.7	72.4 ± 0.5	81.3 ± 0.6	80.6 ± 0.5

1142 Table 20: Scalability of SOLAR on the ViT-G/14 model. Results show top-1 accuracy (%) on full
 1143 datasets.
 1144

Method	# Params	CIFAR-10	CIFAR-100	Food-101	T-ImageNet
LoRA ($r = 4$)	492K	99.4	94.6	91.2	92.8
SOLAR ($r = 4, 8K \rightarrow 4K$)	340K (31% \downarrow)	99.4	94.5	91.2	92.8

1149 Table 21: Comparison of SOLAR with a budget-matched LoRA ($r=2$) baseline on ViT-B. While
 1150 LoRA ($r=2$) has a similar parameter count to the compressed SOLAR adapter, it shows a clear
 1151 performance degradation. SOLAR maintains performance comparable to the original, higher-rank
 1152 LoRA ($r=4$).
 1153

Method	#Params	CIFAR-10	CIFAR-100	Food-101	T-ImageNet
LoRA ($r = 4$)	74K	98.3	90.3	87.6	88.8
LoRA ($r = 2$)	37K	97.1	89.0	85.5	87.4
SOLAR ($r = 4, 4K \rightarrow 1.6K$)	41K	98.3	89.8	87.0	87.9
SOLAR ($r = 4, 4K \rightarrow 0.8K$)	22K	97.0	89.0	85.2	87.4

1160 Table 22: Comparison of SOLAR and simple SVD truncation against standard LoRA adapters on
 1161 multiple vision datasets. The table reports classification accuracy and the corresponding byte footprint
 1162 of the adapter parameters after compression. SOLAR consistently reduces the parameter size while
 1163 preserving or improving performance.
 1164

Method	Byte Footprint	Oxford Pets	SUN397	CUB-200	ImageNet-1K
LoRA ($r = 1$)	74KB	93.0	74.3	84.7	81.5
LoRA ($r = 4$)	297KB	94.2	75.6	86.0	82.8
SVD truncation on LoRA	74KB	92.7	73.3	83.6	80.8
SOLAR on LoRA ($r = 1$)	8KB	92.6	73.9	84.2	81.3
SOLAR on LoRA ($r = 4$)	8KB	93.9	75.0	85.4	82.4

1172
 1173 does not disproportionately harm aggregation performance, even under significant data heterogeneity.
 1174 Our experiment confirms that SOLAR can serve as a post-training, plug-and-play module to
 1175 reduce communication costs in standard FL frameworks without requiring complex changes to the
 1176 aggregation strategy.
 1177

1178 Table 23: Performance of SOLAR on ViT-B under IID and non-IID data distributions in a simulated
 1179 10-client federated learning environment.
 1180

Method	# Params	CIFAR-10 (IID)	CIFAR-10 (non-IID)
LoRA ($r = 4$)	74K	93.7	87.4
SOLAR ($r = 4, 4K \rightarrow 2K$)	51K (31% \downarrow)	93.2	86.7