
Provable Meta-Learning with Low-Rank Adaptations

Jacob L. Block
UT Austin
jblock@utexas.edu

Sundararajan Srinivasan
UT Austin
sundararajans@utexas.edu

Liam Collins
Snap, Inc.
1collins2@snapchat.com

Aryan Mokhtari
UT Austin & Google Research
mokhtari@austin.utexas.edu

Sanjay Shakkottai
UT Austin
sanjay.shakkottai@utexas.edu

Abstract

The power of *foundation models* (FMs) lies in their capacity to learn highly expressive representations that can be adapted to a broad spectrum of tasks. However, these pretrained models require additional training stages to become effective for downstream applications. In the multi-task setting, prior works have shown empirically that specific meta-learning approaches for preparing a model for future adaptation through parameter-efficient fine-tuning (PEFT) can outperform standard retraining methods, but the mechanism of the benefits of meta-learning has been largely unexplored. We introduce a framework for generic PEFT-based meta-learning to learn a model that can easily adapt to unseen tasks. For linear models using LoRA, we show that standard retraining is provably suboptimal for finding an adaptable set of parameters and provide strict performance guarantees for our proposed method. We verify these theoretical insights through experiments on synthetic data as well as real-data vision and language tasks. We observe significant performance benefits using a simple implementation of our proposed meta-learning scheme during retraining relative to the conventional approach.

1 Introduction

Foundation Models (FMs) learn rich representations that are useful for a variety of downstream tasks. The first stage of FM training is referred to as *pretraining*, where a combination of massive public, proprietary, and synthetic sources of data is used to learn a general-purpose model from scratch [1–4]. However, due to the enormous cost of training state-of-the-art models on such large datasets, pretraining is largely infeasible for most. Thus, the most popular and viable way to utilize FMs for specific applications is to adapt an existing pretrained model.

We consider this problem of adapting a pretrained FM to a set of related tasks. We refer to this as *retraining*, where given a number of tasks with many samples, our goal is to recover a model that learns the task structure and can be quickly adapted to future tasks with limited samples. In other works this has been referred to as pre-finetuning [5] or supervised fine-tuning [6]. After retraining, we adapt the model to a new task with few samples in what we denote the *fine-tuning* stage. In this last stage, we typically employ parameter-efficient fine-tuning (PEFT) methods – training heuristics which sacrifice learning expressiveness for improved computational efficiency [7, 8]. Ultimately, the purpose of retraining is to prepare the model for efficient future adaptation, and the effectiveness of a retraining method is measured by the model’s performance on the fine-tuning task.

Standard approaches to retraining involve fitting the model to the aggregation of the different retraining tasks. While this seems reasonable and has been empirically successful [9, 10], it does not leverage knowledge of the downstream fine-tuning procedure to tailor the retrained model to perform well after

such adaptation. Rather, it retrains the model to minimize the average loss across the retraining tasks regardless of the PEFT method to be employed later. Thus, there is no assurance that the recovered solution is indeed adaptable to future unseen tasks relative to other possible retraining solutions.

Meta-learning is a natural framework to address this issue, as it explicitly aims to learn adaptable models, typically in low-resource, few-shot settings using gradient-based adaptations [11, 12]. The success of meta-learning algorithms is largely attributed to their ability to learn useful representations, as even model-agnostic gradient-based meta-learning algorithms like MAML [11] and Reptile [13] have been shown to implicitly learn representations in linear settings [14, 15]. Recent works have shown empirical benefits to retraining using specific PEFT-based meta-learning methods [16–20], but theoretical guarantees showcasing these gains have not been established.

Contributions. We propose a framework for PEFT-based meta-learning during retraining which explicitly accounts for the specific PEFT method that will be used for future fine-tuning. While our framework is general to any PEFT method, we focus on Low-Rank Adaptation (LoRA) [7] for our theoretical results. Specifically, we consider multiple linear regression tasks where each ground truth regressor is a rank- k perturbation of a common matrix $\mathbf{A}^* \in \mathbb{R}^{d \times d}$, where d is the ambient dimension. Given a set of T tasks, our goal is to recover \mathbf{A}^* so that we can easily fine-tune to a new, unseen task by learning a low-rank perturbation using LoRA. With this in mind, we show the following:

- We prove that standard retraining, which does not incorporate any knowledge of the PEFT method for fine-tuning, fails to recover parameters that are low-rank adaptable, as the recovered model is not low-rank away from \mathbf{A}^* and consequently the test model (Proposition 1). As a result, applying low-rank adaptations cannot account for the large rank discrepancy to the test task (Proposition 2). Further, fine-tuning with a very large rank to account for this discrepancy defeats the purpose of PEFT and results in squared prediction error which grows as $\mathcal{O}\left(\frac{d \min\{kT, d\}}{n}\right)$ with high probability, where n is the number of test samples (Remark 2). Thus, standard retraining performs *worse* when given access to more tasks.
- For our meta-learning framework, we guarantee that any minimizer of the meta-learning loss in the infinite sample case is low-rank adaptable to unseen tasks (Theorem 2). Further, we show that if there are at least three retraining tasks, the ground truth parameters are the unique global minima up to orthogonal symmetry (Theorem 3). As a result, LoRA fine-tuning is effective in adapting to the test task and with high probability achieves squared prediction error which grows as $\mathcal{O}\left(\frac{kd}{n}\right)$ (Corollary 4). In contrast to standard retraining, we achieve the optimal rate which does not include any dependence on T .
- We prove in the infinite sample case, every second-order stationary point of our meta-learning loss when applied to two retraining tasks is in fact globally optimal (Theorem 4). In this case there are no spurious local minima of our meta-learning loss and optimality is completely determined by second-order information. Thus, local optimization methods like perturbed gradient descent can efficiently find global minima.

To the author’s knowledge, these are the first results proving that PEFT-based meta-learning outperforms standard retraining methods in any setting. Proofs are presented in Appendix B. To verify our theoretical insights, we compare the performance of the standard retraining and LoRA-based meta-learning objectives in a synthetic multi-output linear regression setting. We show clear improvements using LoRA-based meta-learning for all data parameter settings. Then, we apply an implementation of our general PEFT-based meta-learning framework to a vision task on CIFAR-10 [21] as well as a natural language experiment using the RoBERTa [22] language model on the ConvAI2 [23] dataset with two PEFT schemes: LoRA and last layer fine-tuning. In both cases, we observe that PEFT-based meta-learning outperforms standard retraining.

1.1 Related Work

Meta-learning is a technique for learning models that can be rapidly adapted to unseen tasks by leveraging access to prior tasks during training. For example, Model-Agnostic Meta-Learning (MAML) [11] and Reptile [13] are popular methods that aim to find a model that can be adapted to a new task after a small number of steps of gradient descent on the new task’s loss function.

Prior analyses of gradient-based meta-learning for linear models consider learning a shared feature subspace across tasks [14, 24, 15]. This reduces to a subspace recovery problem, as each task reveals

a linear measurement of this shared space. In contrast, we study tasks that are low-rank perturbations of a shared arbitrary matrix—a novel setting with distinct challenges. For instance, in our setting just three retraining tasks suffice to exactly recover the ground-truth parameters (Theorem 3), whereas subspace-based approaches require the number of tasks to exceed the subspace dimension. These differences motivate new analytical tools and lead to qualitatively different results.

Many meta-learning approaches specific to FMs for incorporating PEFT-based adaptation have been proposed. Hong and Jang [17], Bansal et al. [18], and Gheini et al. [19] applied meta-learning with architecture adaptations that inject task-specific trainable layers within the FM architecture. Hou et al. [16] combined architecture adaptations with parameter perturbations similar to LoRA. They considered a complicated loss that updates the adapters and FM weights over different splits of the data and showed empirical gains over standard retraining and other gradient-based MAML-style algorithms. Aghajanyan et al. [5] proposed a multi-task objective that trains an FM on different tasks simultaneously to encourage learning a universally applicable representation. It forces the FM to learn a shared data representation but allows for task-specific prediction heads. Overall, these works each proposed a meta-learning or multi-task objective and showed empirical gains over standard retraining strategies. However, their experimental results motivate a deeper theoretical exploration of when standard retraining is insufficient relative to meta-learning approaches, how many tasks are needed to learn a rich representation, and how to best adapt to tasks unseen in the training stage.

Lastly, although we focus on LoRA, different PEFT methods have been proposed, including variants of LoRA [25–27] and architecture adaptations [28] among others. Further, recent works have analyzed the theoretical aspects of LoRA in the fine-tuning stage [29, 30], but they explored orthogonal directions to the analysis of LoRA-based meta-learning during retraining. Extended discussion of these prior works is in Appendix A.

Notation. We use bold capital letters for matrices and bold lowercase letters for vectors. $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ refers to the multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. \mathbf{I}_d refers to the $d \times d$ identity matrix. $\|\cdot\|_F$ refers to the Frobenius norm. S_d refers to the set of $d \times d$ symmetric matrices, and S_d^+ is the set of $d \times d$ positive semi-definite matrices. O_d refers to the set of $d \times d$ orthogonal matrices. $[n]$ refers to the set $\{1, \dots, n\}$. For a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, $\text{im}(\mathbf{X})$ and $\text{ker}(\mathbf{X})$ refer to the image and kernel of \mathbf{X} , while $\text{vec}(\mathbf{X}) \in \mathbb{R}^{mn}$ denotes the column-wise vectorization of \mathbf{X} . For subspaces M, N , $\dim(M)$ refers to the dimension of M and $M + N = \{\mathbf{x} + \mathbf{y} | \mathbf{x} \in M, \mathbf{y} \in N\}$. If $M \cap N = \{\mathbf{0}\}$, we write the direct sum $M \oplus N$.

2 Retraining and Fine-Tuning Schemes

We briefly recap the optimization process for standard retraining of an FM across multiple tasks followed by fine-tuning on a downstream task. We then introduce a general framework for PEFT-based meta-learning which adjusts the retraining phase to incorporate insights from fine-tuning.

2.1 Standard Retraining Then Fine-Tuning

Consider a collection of T tasks of interest $\mathcal{T} = \{\mathcal{T}_t\}_{t=1}^T$ where each task \mathcal{T}_t is drawn from task distribution \mathcal{D} and consists of n_t labeled examples $\mathcal{T}_t = \{(\mathbf{x}_{t,j}, \mathbf{y}_{t,j})\}_{j=1}^{n_t}$. Without loss of generality, we assume consistent dimensions across tasks, so $\mathbf{x}_{t,j} \in \mathbb{R}^{d_x}$, $\mathbf{y}_{t,j} \in \mathbb{R}^{d_y}$ for all tasks t and sample indices j . Let $\mathbf{X}_t \in \mathbb{R}^{d_x \times n_t}$ and $\mathbf{Y}_t \in \mathbb{R}^{d_y \times n_t}$ denote the concatenation of the respective input samples and labels from task t , and consider a model $\Phi(\cdot; \mathbf{W}) : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ parameterized by weights \mathbf{W} that maps feature vectors to predicted labels. We abuse notation and write $\Phi(\mathbf{X}_t; \mathbf{W})$ to denote the concatenation of $\Phi(\mathbf{x}_{t,j}; \mathbf{W})$ for $j \in [n_t]$. Typically $\mathbf{W} = (\mathbf{W}_1, \dots, \mathbf{W}_m)$ is a list of matrices where $\mathbf{W}_i \in \mathbb{R}^{d \times d}$ parameterizes the i^{th} layer of a neural network. We assume each \mathbf{W}_i is square for convenience.

Retraining Phase. Given a loss function \mathcal{L} , standard retraining attempts to minimize the aggregated loss over a collection of training tasks [22, 2]. This amounts to solving

$$\hat{\mathbf{W}}_{\text{SR}} = \min_{\mathbf{W}} \sum_{t=1}^T \mathcal{L}(\Phi(\mathbf{X}_t; \mathbf{W}), \mathbf{Y}_t), \quad (1)$$

where SR stands for Standard Retraining. The above optimization problem seeks a set of universal parameters that define a unique mapping function capable of translating inputs to outputs across all tasks involved in the retraining phase. We denote the corresponding model as $\Phi(\cdot; \hat{\mathbf{W}}_{\text{SR}})$.

Fine-Tuning Phase. In the subsequent fine-tuning, PEFT is used to refine the model to fit a downstream task with fewer labeled samples. Formally, consider an unseen task \mathcal{T}_{T+1} drawn from the same task distribution \mathcal{D} . We define PEFT as any method which fits the model to task \mathcal{T}_{T+1} by fixing $\mathbf{W} = \hat{\mathbf{W}}_{\text{SR}}$ in the original parameterization and fine-tuning the mapping $\Phi(\cdot; \hat{\mathbf{W}}_{\text{SR}})$ using additional parameters θ . For example, θ could parameterize trainable perturbations of $\hat{\mathbf{W}}_{\text{SR}}$ or new trainable layers inserted into the architecture of the retrained model [7, 25, 5]. We denote the fine-tuned model as $\Phi_{\text{FT}}(\cdot; \hat{\mathbf{W}}_{\text{SR}}, \theta) : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ and again abuse notation by writing $\Phi_{\text{FT}}(\mathbf{X}_{T+1}; \hat{\mathbf{W}}_{\text{SR}}, \theta)$ to denote the concatenation of $\Phi_{\text{FT}}(\mathbf{x}_{T+1,j}; \hat{\mathbf{W}}_{\text{SR}}, \theta)$ for each $j \in [n_{T+1}]$. During the *fine-tuning stage*, the goal is to find the optimal additional parameters, θ , that minimize the loss for the downstream task \mathcal{T}_{T+1} , solving

$$\min_{\theta} \mathcal{L}(\Phi_{\text{FT}}(\mathbf{X}_{T+1}; \hat{\mathbf{W}}_{\text{SR}}, \theta), \mathbf{Y}_{T+1}). \quad (2)$$

In particular, when LoRA is used for fine-tuning, the model is adapted to task \mathcal{T}_{T+1} by fixing the architecture and the retrained weights $\hat{\mathbf{W}}_{\text{SR}}$ and only training low-rank perturbations for each of the matrices $\hat{\mathbf{W}}_{\text{SR},1}, \dots, \hat{\mathbf{W}}_{\text{SR},m}$. For rank- r adaptations, we parameterize $\theta = ((\mathbf{Q}_1, \mathbf{V}_1), \dots, (\mathbf{Q}_m, \mathbf{V}_m))$, where $\mathbf{Q}_i, \mathbf{V}_i \in \mathbb{R}^{d \times r}$ are the factors of the low-rank adaptation of the i^{th} matrix in $\hat{\mathbf{W}}_{\text{SR}}$. The fine-tuned model is just the original model where the i^{th} weight matrix \mathbf{W}_i is now perturbed to be $\mathbf{W}_i + \mathbf{Q}_i \mathbf{V}_i^{\top}$. For $\mathbf{Q}, \mathbf{V} \in (\mathbb{R}^{d \times r})^m$, define the LoRA loss

$$\mathcal{L}_{\text{LoRA}}(\mathbf{Q}, \mathbf{V}; \mathbf{W}) = \mathcal{L}\left(\Phi\left(\mathbf{X}_{T+1}; (\mathbf{W}_i + \mathbf{Q}_i \mathbf{V}_i^{\top})_{i=1}^m\right), \mathbf{Y}_{T+1}\right). \quad (3)$$

The LoRA fine-tuning optimization problem is then

$$\min_{\mathbf{Q}, \mathbf{V}} \mathcal{L}_{\text{LoRA}}(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{W}}_{\text{SR}}). \quad (4)$$

This pipeline seems reasonable as we first fit the model to the aggregation of the retraining tasks which we hope will promote learning the general structure of the tasks drawn from \mathcal{D} . However, nothing about standard retraining promotes learning an adaptable solution relative to other candidate solutions that fit the retraining tasks. Next, we introduce a general meta-learning framework which explicitly incorporates the adaptation mechanism during retraining.

2.2 PEFT-Based Meta-Learning

Since the ultimate goal of retraining is to perform well on an unseen downstream task, we study a general PEFT-based meta-learning (PEFT-ML) objective that explicitly fits weights and adapter parameters to the training tasks. Rather than training a single model on the aggregation of the retraining tasks, we instead incorporate the adapters during the retraining process and learn adapted models for each task. Let $\theta^{(t)}$ be the set of adapter parameters for the t^{th} training task \mathcal{T}_t . The PEFT-ML objective searches for a single set of base weights $\hat{\mathbf{W}}_{\text{Meta}}$ such that for all $t \in [T]$, the t^{th} adapted model $\Phi_{\text{FT}}(\cdot; \hat{\mathbf{W}}_{\text{Meta}}, \theta^{(t)})$ minimizes the loss over the training task \mathcal{T}_t . More precisely, we define the proposed PEFT-ML objective as

$$\hat{\mathbf{W}}_{\text{Meta}} = \min_{\mathbf{W}} \sum_{t=1}^T \mathcal{L}_t(\mathbf{W}), \quad (5)$$

where $\mathcal{L}_t(\mathbf{W})$ denotes the optimal loss on task t after fine-tuning:

$$\mathcal{L}_t(\mathbf{W}) = \min_{\theta^{(t)}} \mathcal{L}\left(\Phi_{\text{FT}}\left(\mathbf{X}_t; \mathbf{W}, \theta^{(t)}\right), \mathbf{Y}_t\right).$$

When we use LoRA as the adaptation method, we define $\mathbf{Q}^{(t)}, \mathbf{V}^{(t)} \in (\mathbb{R}^{d \times r})^m$ as the list of factors of the low-rank adapter $\mathbf{Q}_i^{(t)} (\mathbf{V}_i^{(t)})^{\top}$ applied to the i^{th} weight matrix for the t^{th} task. Then the inner objective $\mathcal{L}_t(\mathbf{W})$ reduces to

$$\mathcal{L}_t(\mathbf{W}) = \min_{\mathbf{Q}^{(t)}, \mathbf{V}^{(t)}} \mathcal{L}_{\text{LoRA}}(\mathbf{Q}^{(t)}, \mathbf{V}^{(t)}; \mathbf{W}). \quad (6)$$

In this case, we refer to the objective function in (5) as LoRA-ML. This proposed optimization problem is designed to replace the standard retraining objective in (1). After solving (5) we recover base parameters $\hat{\mathbf{W}}_{\text{Meta}}$ that are explicitly designed to be adaptable to downstream tasks drawn from the same distribution as those seen in retraining. To perform finetuning, we then run the exact same minimization in (2) but using retrained weights $\hat{\mathbf{W}}_{\text{Meta}}$ instead of $\hat{\mathbf{W}}_{\text{SR}}$.

3 Main Results

To establish our theoretical results, we consider $T \geq 1$ multi-output linear regression retraining tasks $\{\mathcal{T}_t\}_{t=1}^T$ and one downstream test task \mathcal{T}_{T+1} , where the ground-truth regressor for each task is a low-rank perturbation of a common shared matrix. Precisely, for each $t \in [T+1]$ we assume task \mathcal{T}_t is independently drawn from distribution $\mathcal{D}_{\mathbf{A}^*}$ which is associated with some arbitrary fixed matrix $\mathbf{A}^* \in \mathbb{R}^{d \times d}$ and intrinsic adaptation rank $k \ll d$. Each task $\mathcal{T}_t \sim \mathcal{D}_{\mathbf{A}^*}$ is parameterized by the shared matrix \mathbf{A}^* and a task-specific rank- r matrix \mathbf{R}_t^* such that the samples $(\mathbf{x}_{t,j}, \mathbf{y}_{t,j}) \in \mathbb{R}^d \times \mathbb{R}^d$ from \mathcal{T}_t are related by the noisy linear transformation

$$\mathbf{y}_{t,j} = (\mathbf{A}^* + \mathbf{R}_t^*)\mathbf{x}_{t,j} + \boldsymbol{\epsilon}_{t,j},$$

where $\boldsymbol{\epsilon}_{t,j}$ are independently generated noise terms $\boldsymbol{\epsilon}_{t,j} \sim \mathcal{N}(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I}_d)$. We assume the inputs $\mathbf{x}_{t,j}$ are i.i.d. with zero-mean $\mathbb{E}[\mathbf{x}_{t,j}] = \mathbf{0}$ and covariance $\mathbb{E}[\mathbf{x}_{t,j}\mathbf{x}_{t,j}^\top] = \sigma_x^2 \mathbf{I}_d$. Lastly, we generate the ground truth rank- r adaptations \mathbf{R}_t^* as the symmetric outer product of random Gaussian factors \mathbf{U}_t^* :

$$\mathbf{R}_t^* = \mathbf{U}_t^* \mathbf{U}_t^{*\top} \quad \text{s.t.} \quad \text{vec}(\mathbf{U}_t^*) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{dk}).$$

The above generative model defines the input-output relationships for each task as similar linear models, differing from each other only by a low-rank perturbation \mathbf{R}_t^* . We construct the adapters \mathbf{R}_t^* as symmetric for convenience of analysis, but note that the limitations of standard retraining which we demonstrate in Section 3.1 also hold for general adapters (Appendix C).

Remark 1. *For convenience, we require a mild sense of task diversity and assume that any collection of $r \leq d$ number of columns chosen from the set of all columns of $\{\mathbf{U}_t^*\}_{t=1}^{T+1}$ are linearly independent. We note that our generative model for each \mathbf{U}_t^* ensures this assumption holds with probability 1.*

The learner uses the linear model $\Phi(\mathbf{x}; \mathbf{A}) = \mathbf{A}\mathbf{x}$ for $\mathbf{A} \in \mathbb{R}^{d \times d}$ and retrains on tasks $\mathcal{T}_1, \dots, \mathcal{T}_T$ with the ultimate goal of efficient adaptation to \mathcal{T}_{T+1} using LoRA. The aim is to recover the parameter value $\hat{\mathbf{A}} = \mathbf{A}^*$ in the retraining phase so that the fine-tuned model $\Phi_{\text{FT}}(\mathbf{x}; \hat{\mathbf{A}}, \mathbf{Q}, \mathbf{V}) = (\hat{\mathbf{A}} + \mathbf{Q}\mathbf{V}^\top)\mathbf{x}$ fits the data distribution of any downstream task also drawn from \mathcal{D} for proper rank- k adapter $\mathbf{Q}\mathbf{V}^\top$.

We define the finite-sample loss function for task t as

$$\mathcal{L}_t^{n_t}(\mathbf{A}) = \frac{1}{2n_t} \sum_{j=1}^{n_t} \|\mathbf{y}_{t,j} - \mathbf{A}\mathbf{x}_{t,j}\|_2^2, \quad (7)$$

and we define $\mathcal{L}_t^*(\mathbf{A})$ as the shifted and scaled infinite sample loss:

$$\mathcal{L}_t^*(\mathbf{A}) = \frac{1}{\sigma_x^2} \left(\mathbb{E}_{\mathbf{x}, \mathbf{y}} [\mathcal{L}_t^{n_t}(\mathbf{A})] - \frac{\sigma_\epsilon^2}{2} \right) = \frac{1}{2} \left\| \mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{A} \right\|_F^2. \quad (8)$$

We consider the setting where for the retraining tasks $t \leq T$ we have large n_t , but for the test task n_{T+1} is small. This reflects practical scenarios where we have access to large retraining datasets compared to the low-resource fine-tuning task. Thus, we assume access to the infinite sample loss functions \mathcal{L}_t^* for the retraining tasks $t \leq T$. Then, for ease of notation, define $n = n_{T+1}$ as the number of test task samples. We ultimately aim to use LoRA to fit the finite-sample test task loss \mathcal{L}_{T+1}^n efficiently in n . Given a learned representation $\hat{\mathbf{A}} \in \mathbb{R}^{d \times d}$ from retraining, the fine-tuning problem using LoRA with rank r reduces to

$$\min_{\mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times r}} \mathcal{L}_{T+1}^n(\hat{\mathbf{A}} + \mathbf{Q}\mathbf{V}^\top) \quad (9)$$

Since QV^\top can parameterize any rank- r matrix, (9) is a specific parametrization for what is commonly known as reduced rank regression [31]. It is clear that to even realize the optimal regressor $A^* - \hat{A} + U_{T+1}^* U_{T+1}^{*\top}$ for QV^\top , we need $r \geq \text{rank}(A^* - \hat{A} + U_{T+1}^* U_{T+1}^{*\top})$. Further, results in reduced rank regression and matrix sensing in general reveal the importance of $\text{rank}(A^* + U_{T+1}^* U_{T+1}^{*\top} - \hat{A})$ in terms of the hardness of minimizing $\mathcal{L}_{T+1}^n(\hat{A} + QV^\top)$.

Lemma 1 (Bunea et al., 2011). *Consider $\hat{A} \in \mathbb{R}^{d \times d}$ and let $r = \text{rank}(A^* + U_{T+1}^* U_{T+1}^{*\top} - \hat{A})$. Let $Q^*, V^* \in \mathbb{R}^{d \times r}$ minimize $\mathcal{L}_{T+1}^n(\hat{A} + QV^\top)$ over all rank- r factors $Q, V \in \mathbb{R}^{d \times r}$ and let $X_{T+1} = [x_{T+1,1}, \dots, x_{T+1,n}]$ denote the matrix of test task inputs. Denote the matrix of prediction errors $E = (A^* + U_{T+1}^* U_{T+1}^{*\top})X_{T+1} - (\hat{A} + Q^* V^{*\top})X_{T+1}$. Then $\forall \gamma > 0$,*

$$\mathbb{P}\left(\frac{1}{n} \|E\|_F^2 \leq \frac{24(1+\gamma)^2 \sigma_\epsilon^2 r d}{n} \mid X_{T+1}\right) \geq 1 - e^{-\gamma^2 d} \quad (10)$$

The squared prediction error scales linearly with rd . This matches the information-theoretic lower bound to learn rd number of parameters and is minimax optimal over all rank- r matrices when the singular values of X_{T+1} are uniformly bounded [33]. Thus, a larger rank of $A^* - \hat{A} + U_{T+1}^* U_{T+1}^{*\top}$ inflates the fine-tuning prediction error, as we hope to recover $\hat{A} = A^*$ so that $\text{rank}(A^* - \hat{A} + U_{T+1}^* U_{T+1}^{*\top}) = k$. We next compare the standard retraining (1) and LoRA-ML (6) objectives.

3.1 Negative Results for Standard Retraining then Fine-Tuning

Consider standard retraining then fine-tuning as a candidate for ultimately minimizing (9). The learner first finds a single matrix \hat{A}_{SR} that minimizes the sum of losses $\sum_{t=1}^T \mathcal{L}_t^*$:

$$\hat{A}_{\text{SR}} = \arg \min_A \frac{1}{2} \sum_{t=1}^T \left\| A^* + U_t^* U_t^{*\top} - A \right\|_F^2. \quad (11)$$

Then when given test task \mathcal{T}_{T+1} , the learner solves $\min_{Q, V \in \mathbb{R}^{d \times r}} \mathcal{L}_{T+1}^n(\hat{A}_{\text{SR}} + QV^\top)$. However, this strategy suffers substantial loss in both the retraining and fine-tuning stages. Notice the loss in (11) is convex and quadratic in A , so the first-order optimality condition shows that

$$\hat{A}_{\text{SR}} = A^* + \frac{1}{T} \sum_{t=1}^T U_t^* U_t^{*\top}. \quad (12)$$

Thus, \hat{A}_{SR} recovers A^* added to the average of the retraining ground truth adaptations $U_t^* U_t^{*\top}$. However, \hat{A}_{SR} performs poorly on all of the retraining tasks, as standard retraining is unable to disentangle the common structure A^* from the task-specific adapters $U_t^* U_t^{*\top}$.

Theorem 1. *Let $U^* = (U_1^*, \dots, U_T^*)$. Then,*

$$\mathbb{E}_{U^*} \left[\sum_{t=1}^T \mathcal{L}_t(\hat{A}_{\text{SR}}) \right] = (T-1)kd(d+1) = \Omega(Tkd^2)$$

Thus, \hat{A}_{SR} suffers significant loss on the retraining tasks when averaged over the generation process of ground truth parameters U^* . Further, \hat{A}_{SR} is not low-rank adaptable to the test task. Crucially, the intrinsic dimension of the test task is $\text{rank}(A^* + U_{T+1}^* U_{T+1}^{*\top} - \hat{A}_{\text{SR}}) = \min\{d, k(T+1)\}$, so an adaptation rank of $\min\{d, k(T+1)\}$ is required to even achieve the ground truth test task parameters.

Proposition 1. *If test fine-tuning rank $r < \min\{d, k(T+1)\}$, then $\mathcal{L}_{T+1}^*(Q, V; \hat{A}_{\text{SR}}) > 0$ for all rank- r adapter factors $Q, V \in \mathbb{R}^{d \times r}$.*

Even though the test task parameters are only rank- k away from A^* , standard retraining fails to exploit this structure and inflates the necessary rank to $\min\{d, k(T+1)\}$. Thus, standard retraining actually recovers worse representations as the number of tasks T grows. In this case, failing to fine-tune with large enough rank causes significant loss.

Proposition 2. For a large number of retraining tasks $T \rightarrow \infty$ and test fine-tuning rank r , $\mathcal{L}_{T+1}^*(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{A}}_{SR}) = \Omega((d-r)k^2)$ for all $\mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times r}$.

As the number of retraining tasks grows to infinity, the squared error between the test task recovered parameter and the ground truth is determined by the under-specification of the fine-tuning rank r relative to the ambient dimension d .

The above propositions show the cost of under-specifying the fine-tuning rank relative to the large intrinsic dimension of the test task which results from standard retraining. Conversely, applying the necessarily large fine-tuning rank $r = \min\{d, k(T+1)\}$ both defeats the purpose of *low-rank* adaptation and still incurs large prediction error when fine-tuning with limited samples.

Remark 2. Consider the finite-sample loss (9) using $\hat{\mathbf{A}}_{SR}$ adapted with LoRA using rank $r = \min\{d, k(T+1)\}$. This can achieve optimal population risk but suffers in the finite-sample setting. Using Lemma 1, we can only hope to achieve squared prediction error of order $\mathcal{O}\left(\frac{d \min\{kT, d\}}{n}\right)$ when fine-tuning, which is much larger than the optimal rate $\mathcal{O}\left(\frac{kd}{n}\right)$ if we had in fact recovered the ground truth \mathbf{A}^* during retraining.

Thus, **standard retraining recovers parameters that cannot be efficiently low-rank adapted to new tasks**. In contrast, our analysis of LoRA-ML for retraining shows much improved performance.

3.2 Results for LoRA-Meta-Learning

Consider applying (6) to this problem instance. We introduce low-rank adapters during the retraining phase to model the different training tasks and search for a value of \mathbf{A} such that for all \mathcal{T}_t , the loss \mathcal{L}_t^* after running LoRA on \mathcal{T}_t is minimized. This promotes values of \mathbf{A} that can be easily adapted to unseen tasks downstream. We use the LoRA-ML loss but with symmetric low-rank adapters $\mathbf{U}_t \mathbf{U}_t^\top$ for the t^{th} task \mathcal{T}_t in retraining. We still use asymmetric adapters for fine-tuning on the test task with loss \mathcal{L}_{T+1}^n . The LoRA-ML loss given access to infinite sample task losses \mathcal{L}_t^* is then

$$\mathcal{L}_{\text{Meta}}(\mathbf{A}) = \sum_{t=1}^T \min_{\mathbf{U}_t} \mathcal{L}_t^*(\mathbf{A} + \mathbf{U}_t \mathbf{U}_t^\top). \quad (13)$$

Define the concatenation of each \mathbf{U}_t as $\mathbf{U} = (\mathbf{U}_1, \dots, \mathbf{U}_T) \in (\mathbb{R}^{d \times k})^T$. Then minimizing (13) is equivalent to solving $\min_{\mathbf{A}, \mathbf{U}} \mathcal{L}^*(\mathbf{A}, \mathbf{U})$ where

$$\mathcal{L}^*(\mathbf{A}, \mathbf{U}) = \frac{1}{2} \sum_{t=1}^T \left\| \mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{A} - \mathbf{U}_t \mathbf{U}_t^\top \right\|_F^2. \quad (14)$$

We have seen that standard retraining does not recover an optimal solution, but it is unclear what the global minima of this new objective function are and if they can be easily found. Note that by fixing \mathbf{A} , (14) is T independent symmetric matrix factorization problems, and by fixing \mathbf{U} , (14) is a convex quadratic problem over \mathbf{A} . Despite these well-understood sub-problems, joint minimization over \mathbf{A} and \mathbf{U} presents challenging variable interactions that complicate the analysis. Nevertheless, we employ a careful landscape analysis of (14) to address these questions.

3.2.1 Landscape of Global Minima of (14)

We first show that the objective is well-posed, i.e., minimization of \mathcal{L} leads to an adaptable solution.

Theorem 2. If $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, then $\hat{\mathbf{A}} = \mathbf{A}^* + \mathbf{C}$ where $\text{rank}(\mathbf{C}) \leq 2k$

Any point is a global minimum of (14) if and only if it achieves zero loss. Theorem 2 guarantees that the values of \mathbf{A} that achieve global minimization of (14) are at most rank- $2k$ away from the ground truth parameter \mathbf{A}^* . Then, the remaining intrinsic dimension of the test task is just $3k \ll d$.

Corollary 1. If $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, there exists a rank- $3k$ adapter $\mathbf{Q}\mathbf{V}^\top$ such that $\mathcal{L}_{T+1}^*(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{A}}) = 0$.

Since the sufficient LoRA rank for fine-tuning is just $3k$, we realize a much improved fine-sample prediction error.

Corollary 2. Let $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ and let $\mathbf{Q}^*, \mathbf{V}^* \in \mathbb{R}^{d \times 3k}$ minimize $\mathcal{L}_{T+1}^n(\hat{\mathbf{A}} + \mathbf{Q}\mathbf{V}^\top)$ over all $\mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times 3k}$. Then, $\hat{\mathbf{A}} + \mathbf{Q}^*\mathbf{V}^{*\top}$ satisfies Lemma 1 with $r = 3k$.

Thus, retraining with LoRA-ML leads to squared prediction error on the task task which grows asymptotically as $\mathcal{O}\left(\frac{kd}{n}\right)$. Although the unnecessary factor of T incurred by standard retraining is avoided when using LoRA-ML, the rate still contains an additional factor of 3 over the ideal case when $r = k$ since \mathbf{A}^* is not guaranteed to be recovered exactly. However, this minor discrepancy is mitigated when the number of tasks satisfies $T \geq 3$. In this case, exact recovery of the ground truth parameter \mathbf{A}^* is possible.

Theorem 3. For any number of tasks $T \geq 3$ and ambient dimension $d \geq 3k$, if $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ then $\hat{\mathbf{A}} = \mathbf{A}^*$ and $\mathbf{U}_t\mathbf{U}_t^\top = \mathbf{U}_t^*\mathbf{U}_t^{*\top}$ for all $t \in [T]$

This guarantees that the ground truth parameters are the unique global minimum up to orthogonal symmetry when there are three or more tasks, as long as the ground truth adaptation rank k is relatively small compared to the ambient dimension d . This result is surprising, as most theoretical results for multi-task learning require stricter conditions on the number of tasks T , typically where T is required to be larger than the effective task dimension [34, 14]. However, we establish this uniqueness result for the absolute condition $T \geq 3$. As a result, we only need a rank- k adaptation to realize the test task.

Corollary 3. For $T \geq 3$, if $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, then $\exists \mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times k}$ such that $\mathcal{L}_{T+1}^*(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{A}}) = 0$.

We then achieve the desired fine-sample prediction error.

Corollary 4. For $T \geq 3$, let $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ and let $\mathbf{Q}^*, \mathbf{V}^* \in \mathbb{R}^{d \times k}$ minimize $\mathcal{L}_{T+1}^n(\hat{\mathbf{A}} + \mathbf{Q}\mathbf{V}^\top)$ over all $\mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times k}$. Then, $\hat{\mathbf{A}} + \mathbf{Q}^*\mathbf{V}^{*\top}$ satisfies Lemma 1 with $r = k$.

Note that the condition $T \geq 3$ is necessary to establish Theorem 3, as if there are only two tasks we can construct ground truth parameters such that the induced loss \mathcal{L}^* has infinite solutions. See Appendix F.1 for an example.

Summary. These results show that all global minima of the LoRA-ML objective are low-rank adaptable to the downstream task and achieve finite-sample test task prediction error which grows as $\mathcal{O}\left(\frac{kd}{n}\right)$. Crucially, this avoids the factor of T incurred by standard retraining. Further, if $T \geq 3$, minimizing the LoRA-ML objective guarantees recovery of the ground truth parameters.

3.2.2 Algorithms for Minimizing (14)

As shown above, minimizing the LoRA-ML objective (14) leads to recovery of the ground truth parameters, with a small rank- $2k$ error term when $T = 2$. We prove that this minimization problem can always be solved by local optimization methods when there are two retraining tasks.

Theorem 4. If $T = 2$ and the ambient dimension $d \geq 2k$, then $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ if and only if $(\hat{\mathbf{A}}, \hat{\mathbf{U}})$ is a second order stationary point (SOSP) of \mathcal{L}^* .

When $T = 2$, local optimization algorithms for finding SOSPs, such as perturbed gradient descent and cubic-regularized Newton method, can efficiently minimize the meta-learning objective. Surprisingly, when there are three or more tasks, numerical experiments (see Appendix F.2) show that adversarially picking \mathbf{U}_t^* can result in specific instantiations of (14) with spurious local minima. In the next section, we perform extensive numerical experiments for various values of T which show that these spurious minima are almost never found in practice and vanilla gradient descent is sufficient to minimize (14).

4 Experiments

We test our framework across three settings. We first consider synthetic data for linear models using LoRA fine-tuning to validate our theory in Section 3. Next, we conduct real-data experiments on vision and language tasks, comparing our general PEFT-ML retraining approach to standard retraining and Reptile [13], a gradient-based meta-learning method. To highlight the flexibility of our framework, we consider two different fine-tuning schemes: LoRA and fine-tuning just the last layer.

Synthetic Data. We first test our model on synthetic regression tasks, relaxing the assumptions from our theoretical results. We consider data $\mathbf{y}_{t,j} = (\mathbf{A}^* + \mathbf{R}_t^*)\mathbf{x}_{t,j} + \epsilon_{t,j}$ for task t and $j \in [n_t]$, where




Figure 1: Linear model fine-tuning performance varying number of retraining tasks T (left) and number of fine-tuning samples n (right) for LoRA-ML (ours) and standard retraining (SR).

Table 1: Mean test accuracies and standard errors for PEFT-ML and standard retraining methods using LoRA-based and last-layer-based fine-tuning for adapting to a subset of CIFAR-10 classes.

LoRA Fine-Tuning			Last-Layer Fine-Tuning		
Method	Mean	Std. Err.	Method	Mean	Std. Err.
LoRA-ML	86.09	0.35	Last-Layer-ML	83.90	0.17
SR+LoRA	85.29	0.17	SR+Last-Layer	81.16	0.22
Reptile+LoRA	81.30	0.42	Reptile+Last-Layer	72.55	0.39

$\mathbf{x}_{t,j}$, $\epsilon_{t,j}$, and \mathbf{R}_t^* are generated just as in Section 3. \mathbf{A}^* is constructed by sampling each entry as an i.i.d. $\mathcal{N}(0, 1)$ random variable. We define parameters N, n such that the number of samples per retraining task $n_t = N$ for all $t \leq T$, and the number of test task samples $n_{T+1} = n$. Setting \mathcal{L} to be the mean squared error loss, we run gradient descent on the standard retraining (1) and the LoRA-ML (5) objectives. After recovering $\hat{\mathbf{A}}$ during retraining, we apply the low-rank adaptation $\mathbf{Q}\mathbf{V}^\top$ when fine-tuning to the test task. When $T = 2$, we use a rank- $3k$ adaptation during fine-tuning to account for the inexact recovery explained in Theorem 2, and otherwise use a rank- k adaptation.

In each experiment we vary one hyperparameter from a fixed set of values and plot the prediction error between the recovered model and the ground truth model $\frac{1}{n} \sum_j \|\mathbf{y}_{T+1,j} - (\hat{\mathbf{A}} + \mathbf{Q}\mathbf{V}^\top)\mathbf{x}_{T+1,j}\|_2^2$. Results are averaged over 5 trials, with the shaded region showing the full range of values. Figure 1 shows that LoRA-ML retraining significantly outperforms standard retraining (SR) for all data settings. For $T > 2$, we see that applying gradient descent to the LoRA-ML objective is sufficient to achieve global minimization and recover an adaptable solution. Thus, even though there may exist spurious local minimizers, we do not encounter them in practice. Further, the fine-tuning performance after standard retraining worsens with larger T . This is supported by our theory in Section 3.1 which shows that as T increases, standard retraining recovers worse solutions that leave a larger intrinsic dimension for the fine-tuning stage. See Appendix D for hyperparameter details and further ablations.

Vision Experiments. We use CIFAR-10 [21], and define $T = 4$ binary retraining tasks involving classification between consecutive class labels. Specifically, task 1 classifies between classes 1 and 2, task 2 between classes 3 and 4, etc. The test task is binary classification between classes 9 and 10. We evaluate our PEFT-ML framework using a model based on the MLP-Mixer architecture [35]. We compare two different adaptation methods: (i) LoRA with rank 1 adapters, and (ii) fine-tuning only the last layer. For both PEFT strategies, we evaluate three retraining strategies: (a) PEFT-ML for each PEFT method (LoRA-ML, Last-Layer-ML), (b) standard retraining, and (c) Reptile, a popular gradient-based meta-learning method. We report mean results along with standard errors across 5 trials in Table 1. Across both fine-tuning methods, PEFT-ML retraining consistently yields the best performance. Interestingly, standard retraining outperforms Reptile, as Reptile assumes gradient-based adaptation which is misaligned with the PEFT methods used to adapt at test-time.

Language Experiments. We test our PEFT-ML framework using the ConvAI2 dataset with the RoBERTa-base model [22]. ConvAI2 comprises conversations between two personas, where each

Table 2: Mean accuracies \pm standard error for 10 test tasks using different retraining and fine-tuning method combinations on ConvAI2.

Algorithm	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	Avg
LoRA-ML	57 \pm 2	41 \pm 4	51 \pm 4	50 \pm 4	51 \pm 4	30 \pm 2	66 \pm 5	47 \pm 4	43 \pm 4	38 \pm 2	47.4 \pm 1.9
SR+LoRA	59 \pm 5	31 \pm 4	50 \pm 7	40 \pm 4	24 \pm 5	20 \pm 2	41 \pm 10	36 \pm 2	23 \pm 5	26 \pm 6	35.0 \pm 4.1
Reptile+LoRA	45 \pm 7	29 \pm 5	35 \pm 6	36 \pm 2	19 \pm 6	21 \pm 3	28 \pm 10	29 \pm 7	21 \pm 5	21 \pm 7	28.4 \pm 5.0
Last-Layer-ML	55 \pm 2	27 \pm 5	51 \pm 5	44 \pm 2	36 \pm 4	25 \pm 4	55 \pm 4	43 \pm 5	33 \pm 4	34 \pm 4	40.2 \pm 1.5
SR+Last-Layer	41 \pm 4	9 \pm 5	29 \pm 3	36 \pm 5	28 \pm 8	15 \pm 4	31 \pm 8	24 \pm 5	15 \pm 6	17 \pm 6	24.5 \pm 4.2
Reptile+Last-Layer	35 \pm 6	13 \pm 4	18 \pm 3	29 \pm 3	19 \pm 9	18 \pm 4	15 \pm 9	17 \pm 5	15 \pm 6	16 \pm 4	19.4 \pm 4.3

Table 3: Mean accuracies (averaged over all test tasks) \pm standard error for different retraining methods at varying LoRA ranks during fine-tuning.

Algorithm	Rank 1	Rank 4	Rank 8	Rank 16
LoRA-ML	44.7 \pm 0.8	48.6 \pm 1.6	47.4 \pm 1.9	48.2 \pm 1.5
SR+LoRA	36.3 \pm 4.1	35.5 \pm 4.3	35.0 \pm 4.1	37.1 \pm 4.1
Reptile+LoRA	26.6 \pm 5.8	27.7 \pm 5.3	28.4 \pm 5.0	27.8 \pm 5.9

persona is associated with a list of facts that informs their responses. We model learning the dialogue continuations of each persona as a different task, where we aim to select the correct continuation from a set of choices given the conversation history. We retrain with $T = 10$ tasks and then fine-tune the model from the best-performing epoch to each of the 10 test tasks. We run 5 trials and report the mean accuracy and standard error on the held-out data for each test task in Table 2, using the same setup and method naming conventions as in the vision experiments. PEFT-based meta-learning performs best for both adaptation methods.

Table 2 reports results where LoRA-ML retraining and all fine-tuning uses a LoRA rank of 8. To assess robustness to the choice of rank, Table 3 reports performance across different values. LoRA-ML consistently outperforms the baselines, and while choosing a rank larger than 1 clearly improves the performance of LoRA-ML, ranks 4, 8, and 16 yield similar results. We note that LoRA-ML is always retrained with the same rank used for fine-tuning. Additional details for both real-data experiments are provided in Appendix E.

5 Conclusion

We presented PEFT-ML for retraining an FM on a collection of tasks to prepare the model for subsequent downstream fine-tuning. We theoretically demonstrated strict performance gaps between standard retraining and the PEFT-ML objective using LoRA (LoRA-ML). Empirically, retraining with PEFT-ML outperformed standard retraining for adapting to unseen downstream tasks using LoRA and last-layer fine-tuning. Future research includes extending our theoretical analysis to more general adapters and more complex model architectures, such as transformers.

Acknowledgments

This work was supported in part by NSF Grants 2019844, 2107037, and 2112471, ONR Grant N00014-19-1-2566, the Machine Learning Lab (MLL) at UT Austin, the NSF AI Institute for Foundations of Machine Learning (IFML), and Qualcomm through the Wireless Networking and Communications Group (WNCG) Industrial Affiliates Program. We are grateful for computing support on the Vista GPU Cluster through the Center for Generative AI (CGAI) and the Texas Advanced Computing Center (TACC) at the University of Texas at Austin.

References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [3] Marah I Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Hassan Awadalla, Nguyen Bach, et al. Phi-3 technical report: A highly capable language model locally on your phone. Technical Report MSR-TR-2024-12, Microsoft, 2024.
- [4] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [5] Armen Aghajanyan, Anchit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. Muppet: Massive multi-task representations with pre-finetuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2021.
- [6] Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. How abilities in large language models are affected by supervised fine-tuning data composition. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 177–198, 2024.
- [7] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.
- [8] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, 2021.
- [9] Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. Unifiedqa: Crossing format boundaries with a single qa system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, 2020.
- [10] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- [12] Yoonho Lee and Seungjin Choi. Gradient-based meta-learning with learned layerwise metric and subspace. In *International conference on machine learning*, pages 2927–2936. PMLR, 2018.
- [13] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [14] Liam Collins, Aryan Mokhtari, Sewoong Oh, and Sanjay Shakkottai. MAML and ANIL provably learn representations. In *International conference on machine learning*, pages 4238–4310. PMLR, 2022.

- [15] Nikunj Saunshi, Yi Zhang, Mikhail Khodak, and Sanjeev Arora. A sample complexity separation between non-convex and convex meta-learning. In *International Conference on Machine Learning*, pages 8512–8521. PMLR, 2020.
- [16] Zejiang Hou, Julian Salazar, and George Polovets. Meta-learning the difference: Preparing large language models for efficient adaptation. *Transactions of the Association for Computational Linguistics*, 10:1249–1265, 2022.
- [17] SK Hong and Tae Young Jang. AMAL: Meta knowledge-driven few-shot adapter learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10381–10389, 2022.
- [18] Trapit Bansal, Salaheddin Alzubi, Tong Wang, Jay-Yoon Lee, and Andrew McCallum. Meta-adapters: Parameter efficient few-shot fine-tuning through meta-learning. In *International Conference on Automated Machine Learning*, pages 19–1. PMLR, 2022.
- [19] Mozhddeh Gheini, Xuezhe Ma, and Jonathan May. Know where you’re going: Meta-learning for parameter-efficient fine-tuning. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 11602–11612, 2023.
- [20] Nathan Hu, Eric Mitchell, Christopher D Manning, and Chelsea Finn. Meta-learning online adaptation of language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4418–4432, 2023.
- [21] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [22] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [23] Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander H. Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, Shrimai Prabhumoye, Alan W. Black, Alexander I. Rudnicky, Jason Williams, Joelle Pineau, Mikhail S. Burtsev, and Jason Weston. The second conversational intelligence challenge (ConvAI2). *ArXiv*, abs/1902.00098, 2019.
- [24] Kiran K Thekumparampil, Prateek Jain, Praneeth Netrapalli, and Sewoong Oh. Statistically and computationally efficient linear meta-representation learning. *Advances in Neural Information Processing Systems*, 34:18487–18500, 2021.
- [25] Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. In *Forty-first International Conference on Machine Learning*, 2024.
- [26] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. *Advances in neural information processing systems*, 36:10088–10115, 2023.
- [27] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive budget allocation for parameter-efficient fine-tuning. In *The Eleventh International Conference on Learning Representations*, 2023.
- [28] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR, 6 2019.
- [29] Uijeong Jang, Jason D Lee, and Ernest K Ryu. LoRA training in the NTK regime has no spurious local minima. In *Proceedings of the 41st International Conference on Machine Learning*, pages 21306–21328, 2024.

- [30] Yuchen Zeng and Kangwook Lee. The expressive power of low-rank adaptation. In *The Twelfth International Conference on Learning Representations*, 2024.
- [31] Alan Julian Izenman. Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis*, 5(2):248–264, 1975. ISSN 0047-259X. doi: [https://doi.org/10.1016/0047-259X\(75\)90042-1](https://doi.org/10.1016/0047-259X(75)90042-1).
- [32] Florentina Bunea, Yiyuan She, and Marten H. Wegkamp. Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics*, 39(2):1282–1309, 2011. ISSN 00905364, 21688966.
- [33] Angelika Rohde and Alexandre B. Tsybakov. Estimation of high-dimensional low-rank matrices. *The Annals of Statistics*, 39(2):887–930, 2011. ISSN 00905364, 21688966.
- [34] Simon Shaolei Du, Wei Hu, Sham M. Kakade, Jason D. Lee, and Qi Lei. Few-shot learning via learning the representation, provably. In *International Conference on Learning Representations*, 2021.
- [35] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 24261–24272, 2021.
- [36] Yuqian Zhang, Qing Qu, and John Wright. From symmetry to geometry: Tractable nonconvex problems. *arXiv preprint arXiv:2007.06753*, 2020.

A Related Work on LoRA-Style PEFT

There is a vast amount of work in developing PEFT methods for FMs. The LoRA algorithm [7] has established itself as a popular and successful PEFT strategy and has inspired various extensions such as QLoRA, DoRA, and others [26, 25, 27]. These algorithms are heuristics for mimicking the full finetuning of an FM to a specific downstream task and have proven to be empirically successful in various settings. However, there is a lack of theoretical analysis on the adaptability of PFMs under LoRA-style adaptations, the ability to efficiently optimize LoRA-style objectives, and the kinds of solutions they recover. Some recent works have attempted to analyze different parts of these theoretical questions.

Convergence of LoRA. [29] analyzes the optimization landscape for LoRA for the Neural Tangent Kernel regime. The authors show that LoRA finetuning converges in this setting as they prove that the objective function satisfies a strict saddle property, ensuring that there are no spurious local minima. However, this focuses on the actual ability of LoRA to converge to the optimal low-rank adapter given an FM, and does not consider the adaptability of the FM in the first place.

Expressivity of LoRA. [30] derives the expressive power of LoRA as a function of model depth. This work shows that under some mild conditions, fully connected and transformer networks when respectively adapted with LoRA can closely approximate arbitrary smaller networks. They quantify the required LoRA rank to achieve this approximation as well as the resulting approximation error.

B Proofs

B.1 Proof of Theorem 1

By Equation (12) we have that $\hat{\mathbf{A}}_{\text{SR}} = \mathbf{A}^* + \frac{1}{T} \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top}$. In the following the expectation is always taken over $\mathbf{U}^* = (\mathbf{U}_1^*, \dots, \mathbf{U}_T^*)$, where $\mathbf{U}_t^* \in \mathbb{R}^{d \times k}$ satisfies $\text{vec}(\mathbf{U}_t^*) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{dk})$. Then,

$$\begin{aligned}
\mathbb{E}_{\mathbf{U}^*} \left[\sum_{t=1}^T \mathcal{L}_t(\hat{\mathbf{A}}_{\text{SR}}) \right] &= \sum_{t=1}^T \mathbb{E} \left[\left\| \mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{A}^* - \frac{1}{T} \sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} \right\|_F^2 \right] \\
&= \sum_{t=1}^T \mathbb{E} \left[\left\| \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \frac{1}{T} \sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} \right\|_F^2 \right] \\
&= \sum_{t=1}^T \mathbb{E} \left[\left\| \mathbf{U}_t^* \mathbf{U}_t^{*\top} - k\mathbf{I} - \frac{1}{T} \sum_{s=1}^T (\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}) \right\|_F^2 \right] \\
&= \sum_{t=1}^T \mathbb{E} \left[\left\| \mathbf{U}_t^* \mathbf{U}_t^{*\top} - k\mathbf{I} \right\|_F^2 + \left\| \frac{1}{T} \sum_{s=1}^T (\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}) \right\|_F^2 - \right. \\
&\quad \left. \frac{2}{T} \text{tr} \left\{ (\mathbf{U}_t^* \mathbf{U}_t^{*\top} - k\mathbf{I}) \left(\sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I} \right) \right\} \right] \\
&= T \mathbb{E} \left[\left\| \mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I} \right\|_F^2 \right] + \frac{1}{T} \mathbb{E} \left[\left\| \sum_{s=1}^T (\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}) \right\|_F^2 \right] \\
&\quad - 2 \mathbb{E} \left[\text{tr} \left\{ (\mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I}) \left(\sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I} \right) \right\} \right],
\end{aligned}$$

where the last equality follows from the fact that each $\mathbf{U}_t^* \mathbf{U}_t^{*\top}$ are i.i.d.

Note that the second term $\frac{1}{T} \mathbb{E} \left[\left\| \sum_{s=1}^T (\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}) \right\|_F^2 \right]$ is the total variance of the sum of i.i.d. matrices $\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}$, so it is equal to the sum of the individual total variances:

$$\frac{1}{T} \mathbb{E} \left[\left\| \sum_{s=1}^T (\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}) \right\|_F^2 \right] = \frac{1}{T} \sum_{s=1}^T \mathbb{E} \left[\left\| (\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}) \right\|_F^2 \right] = \mathbb{E} \left[\left\| \mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I} \right\|_F^2 \right].$$

Further, the third term $-2\mathbb{E} \left[\text{tr} \left\{ (\mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I}) \left(\sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I} \right) \right\} \right]$ is the sum of total covariances between zero mean random matrices $\mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I}$ and each $\mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}$. Since each \mathbf{U}_s^* is drawn independently, we only pick up the first term of $\sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I}$:

$$\begin{aligned} \mathbb{E} \left[\text{tr} \left\{ (\mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I}) \left(\sum_{s=1}^T \mathbf{U}_s^* \mathbf{U}_s^{*\top} - k\mathbf{I} \right) \right\} \right] &= \mathbb{E} \left[\text{tr} \left\{ (\mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I}) (\mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I}) \right\} \right] \\ &= \mathbb{E} \left[\left\| \mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I} \right\|_F^2 \right]. \end{aligned}$$

Combining the above simplifications gives that

$$\mathbb{E}_{\mathbf{U}^*} \left[\sum_{t=1}^T \mathcal{L}_t(\hat{\mathbf{A}}_{\text{SR}}) \right] = (T-1) \mathbb{E} \left[\left\| \mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I} \right\|_F^2 \right].$$

Then using the fact that each $\mathbf{U}_t^* \mathbf{U}_t^{*\top}$ is an independent sample of a $d \times d$ Wishart distribution with scale matrix \mathbf{I} and k degrees of freedom, we have that

$$\mathbb{E} \left[\left\| \mathbf{U}_1^* \mathbf{U}_1^{*\top} - k\mathbf{I} \right\|_F^2 \right] = kd(d+1)$$

B.2 Proof of Propositions 1, 2

Recall $\hat{\mathbf{A}}_{\text{SR}} = \mathbf{A}^* + \frac{1}{T} \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top}$. Then for any $\mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times r}$,

$$\begin{aligned} \mathcal{L}_{T+1}^*(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{A}}_{\text{SR}}) &= \frac{1}{2} \left\| \mathbf{A}^* + \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - \hat{\mathbf{A}}_{\text{SR}} - \mathbf{Q}\mathbf{V}^\top \right\|_F^2 \\ &= \frac{1}{2} \left\| \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{Q}\mathbf{V}^\top \right\|_F^2 \end{aligned}$$

By Remark 1 we have that $\text{rank} \left(\mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top} \right) = \min\{d, k(T+1)\}$. Then Proposition 1 follows from that fact that $\text{rank}(\mathbf{Q}\mathbf{V}^\top) \leq r$.

Further, as $T \rightarrow \infty$, the strong law of large numbers implies that $\frac{1}{T} \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top} \rightarrow \mathbb{E} \left[\mathbf{U}_t^* \mathbf{U}_t^{*\top} \right] = k\mathbf{I}$. Thus for large T ,

$$\left\| \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - \frac{1}{T} \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{Q}\mathbf{V}^\top \right\|_F^2 \rightarrow \left\| \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - k\mathbf{I} - \mathbf{Q}\mathbf{V}^\top \right\|_F^2 \quad (15)$$

Using classic low-rank matrix factorization results, the $\mathbf{Q}^* \mathbf{V}^{*\top}$ that minimizes $\left\| \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - k\mathbf{I} - \mathbf{Q}\mathbf{V}^\top \right\|_F^2$ will exactly capture the r eigenvectors of $\mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - k\mathbf{I}$ with largest magnitude eigenvalue. But, $\mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - k\mathbf{I}$ has $d-k$ eigenvalues of magnitude k , so $\mathbf{Q}^* \mathbf{V}^{*\top}$ can only capture r of them. Thus, $\left\| \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - k\mathbf{I} - \mathbf{Q}^* \mathbf{V}^{*\top} \right\|_F^2 \geq (d-k-r)k^2$. Since $\mathbf{Q}^* \mathbf{V}^{*\top}$ minimized this quantity, we have that

$$\mathcal{L}_{T+1}^*(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{A}}_{\text{SR}}) \geq (d-k-r)k^2 \quad \forall \mathbf{Q}, \mathbf{V} \in \mathbb{R}^{d \times r}$$

Thus, $\mathcal{L}_{T+1}(\mathbf{Q}, \mathbf{V}; \hat{\mathbf{A}}_{\text{SR}})$ scales as $(d-k-r)k^2 \approx (d-r)k^2$ since $k \ll d$.

B.3 Proof of Theorem 2

Since $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ and $\mathcal{L}^* \geq 0$ we must have that $\nabla_{\mathbf{A}} \mathcal{L}^* = 0$.

Thus, $\hat{\mathbf{A}} = \mathbf{A}^* - \frac{1}{T} \sum_{j=1}^T (\hat{\mathbf{U}}_j \hat{\mathbf{U}}_j^\top - \mathbf{U}_j^* \mathbf{U}_j^{*\top})$. Plugging this into \mathcal{L}^* gives

$$\begin{aligned} 0 = \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) &= \frac{1}{2} \sum_{t=1}^T \left\| \mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \left(\mathbf{A}^* - \frac{1}{T} \sum_{s=1}^T (\hat{\mathbf{U}}_s \hat{\mathbf{U}}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top}) \right) - \mathbf{U}_t \mathbf{U}_t^\top \right\|_F^2 \\ &= \frac{1}{2} \sum_{t=1}^T \left\| \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{U}_t \mathbf{U}_t^\top - \frac{1}{T} \sum_{s=1}^T (\hat{\mathbf{U}}_s \hat{\mathbf{U}}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top}) \right\|_F^2. \end{aligned}$$

Thus each term of the summation is zero, so for all $t, s \in [T]$,

$$\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top} = \hat{\mathbf{U}}_s \hat{\mathbf{U}}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top}.$$

Combining these results gives that

$$\begin{aligned} \hat{\mathbf{A}} &= \mathbf{A}^* - \frac{1}{T} \sum_{s=1}^T (\hat{\mathbf{U}}_s \hat{\mathbf{U}}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top}) \\ &= \mathbf{A}^* - (\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top - \mathbf{U}_1^* \mathbf{U}_1^{*\top}) \end{aligned}$$

Let $\mathbf{C} = -\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top + \mathbf{U}_1^* \mathbf{U}_1^{*\top}$.

Then $\hat{\mathbf{A}} = \mathbf{A}^* + \mathbf{C}$ and $\text{rank}(\mathbf{C}) \leq \text{rank}(\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top) + \text{rank}(\mathbf{U}_1^* \mathbf{U}_1^{*\top}) \leq 2k$.

Note the the effective remaining test-task dimension is

$$\begin{aligned} \text{rank}(\mathbf{A}^* + \mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - \mathbf{A}^* - \mathbf{C}) &= \text{rank}(\mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top} - \mathbf{C}) \\ &\leq \text{rank}(\mathbf{U}_{T+1}^* \mathbf{U}_{T+1}^{*\top}) + \text{rank}(\mathbf{C}) \\ &\leq 3k \end{aligned}$$

B.4 Proof of Theorem 3

Proof. Since $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, we have that for all $t, s \in [T]$,

$$\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top} = \hat{\mathbf{U}}_s \hat{\mathbf{U}}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top} \quad (16)$$

Applying this to the first three tasks and rearranging gives that

$$\mathbf{U}_1^* \mathbf{U}_1^{*\top} = \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top + \mathbf{U}_2^* \mathbf{U}_2^{*\top} - \hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top \quad (17)$$

$$= \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top + \mathbf{U}_3^* \mathbf{U}_3^{*\top} - \hat{\mathbf{U}}_3 \hat{\mathbf{U}}_3^\top. \quad (18)$$

We first show that $\text{im}(\hat{\mathbf{U}}_1) = \text{im}(\mathbf{U}_1^*)$.

Since $\mathbf{U}_1^* \mathbf{U}_1^{*\top} \succcurlyeq 0$, we must have that $\text{im}(\hat{\mathbf{U}}_2) \subseteq \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_2^*)$ and $\text{im}(\hat{\mathbf{U}}_3) \subseteq \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_3^*)$, as otherwise there would exist a vector on $\ker(\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top + \mathbf{U}_2^* \mathbf{U}_2^{*\top}) \cap \ker(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top)^\perp$ whose existence contradicts the positive semi-definiteness of $\mathbf{U}_1^* \mathbf{U}_1^{*\top}$.

Thus,

$$\text{im}(\mathbf{U}_1^*) \subseteq \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_2^*) \quad (19)$$

$$\text{im}(\mathbf{U}_1^*) \subseteq \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_3^*) \quad (20)$$

Using that fact that for subspaces $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$, $\mathbf{X} \subseteq \mathbf{Y} \implies \mathbf{X} + \mathbf{Z} \subseteq \mathbf{Y} + \mathbf{Z}$, we can add $\text{im}(\mathbf{U}_2^*)$ and $\text{im}(\mathbf{U}_3^*)$ to both sides of 19 and 20 respectively. This gives:

$$\text{im}(\mathbf{U}_1^*) \oplus \text{im}(\mathbf{U}_2^*) \subseteq \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_2^*) \quad (21)$$

$$\text{im}(\mathbf{U}_1^*) \oplus \text{im}(\mathbf{U}_3^*) \subseteq \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_3^*). \quad (22)$$

For $t \in \{2, 3\}$, we clearly have that $\dim(\text{im}(\hat{\mathbf{U}}_1) + \text{im}(\mathbf{U}_t^*)) \leq \dim \text{im}(\hat{\mathbf{U}}_1) + \dim \text{im}(\mathbf{U}_t^*) \leq 2k$, and $\dim(\text{im}(\mathbf{U}_1^*) + \text{im}(\mathbf{U}_t^*)) = 2k$. Thus,

$$(\text{im}(\mathbf{U}_1^*) \oplus \text{im}(\mathbf{U}_2^*)) = (\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_2^*)) \quad (23)$$

$$(\text{im}(\mathbf{U}_1^*) \oplus \text{im}(\mathbf{U}_3^*)) = (\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_3^*)) \quad (24)$$

Lemma 2. $([\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_2^*)] \cap [\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_3^*)]) = \text{im}(\hat{\mathbf{U}}_1)$

Proof. Clearly, $\text{im}(\hat{\mathbf{U}}_1) \subseteq ([\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_2^*)] \cap [\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_3^*)])$. To show the converse, consider $\mathbf{x} \in ([\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_2^*)] \cap [\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_3^*)])$.

By assumption there exists some $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d} \in \mathbb{R}^k$ such that

$$\mathbf{x} = \hat{\mathbf{U}}_1 \mathbf{a} + \mathbf{U}_2^* \mathbf{b} \quad (25)$$

$$= \hat{\mathbf{U}}_1 \mathbf{c} + \mathbf{U}_3^* \mathbf{d} \quad (26)$$

Thus,

$$\hat{\mathbf{U}}_1(\mathbf{a} - \mathbf{c}) + \mathbf{U}_2^* \mathbf{b} - \mathbf{U}_3^* \mathbf{d} = \mathbf{0}. \quad (27)$$

By Equation 23, we can write

$$\begin{aligned} \text{im}(\mathbf{U}_2^*) &= ([\text{im}(\mathbf{U}_1^*) \oplus \text{im}(\mathbf{U}_2^*)] \cap [\text{im}(\mathbf{U}_2^*) \oplus \text{im}(\mathbf{U}_3^*)]) \\ &= ([\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_2^*)] \cap [\text{im}(\mathbf{U}_2^*) \oplus \text{im}(\mathbf{U}_3^*)]) \end{aligned}$$

Thus, $\text{im}(\hat{\mathbf{U}}_1) \cap [\text{im}(\mathbf{U}_2^*) \oplus \text{im}(\mathbf{U}_3^*)] \subseteq \text{im}(\hat{\mathbf{U}}_1) \cap \text{im}(\mathbf{U}_2^*) = \{\mathbf{0}\}$, so

$$\text{im}(\hat{\mathbf{U}}_1) \cap [\text{im}(\mathbf{U}_2^*) \oplus \text{im}(\mathbf{U}_3^*)] = \{\mathbf{0}\} \quad (28)$$

Applying Equation (28) to Equation (27) implies that $\mathbf{a} = \mathbf{c}$ and $\mathbf{b} = \mathbf{d} = \mathbf{0}$. Thus $\mathbf{x} = \hat{\mathbf{U}}_1 \mathbf{a} \in \text{im}(\hat{\mathbf{U}}_1)$, so $([\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_2^*)] \cap [\text{im}(\hat{\mathbf{U}}_1) \oplus \text{im}(\mathbf{U}_3^*)]) \subseteq \text{im}(\hat{\mathbf{U}}_1)$. \square

Then Equations (19) and (20) combined with Lemma (2) implies that $\text{im}(\mathbf{U}_1^*) \subseteq \text{im}(\hat{\mathbf{U}}_1)$ but $\dim(\text{im}(\mathbf{U}_1^*)) = \dim(\text{im}(\hat{\mathbf{U}}_1)) = k$, so $\text{im}(\mathbf{U}_1^*) = \text{im}(\hat{\mathbf{U}}_1)$.

Since the initial assumptions about $\hat{\mathbf{U}}_1$ and \mathbf{U}_1^* analogously hold for the corresponding matrices for tasks 2 and 3, by the exact same argument we can show that

$$\text{im}(\mathbf{U}_t^*) = \text{im}(\hat{\mathbf{U}}_t) \quad \forall t \in [T]. \quad (29)$$

Then by equation (16), $\text{im}(\mathbf{U}_1^*) \supseteq \text{im}(\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top - \mathbf{U}_1^* \mathbf{U}_1^{*\top}) = \text{im}(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \mathbf{U}_2^* \mathbf{U}_2^{*\top}) \subseteq \text{im}(\mathbf{U}_2^*)$. Thus,

$$\begin{aligned} \text{im}(\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top - \mathbf{U}_1^* \mathbf{U}_1^{*\top}) &\subseteq \text{im}(\mathbf{U}_1^*) \cap \text{im}(\mathbf{U}_2^*) \\ &= \{\mathbf{0}\}. \end{aligned}$$

Thus $\hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top = \mathbf{U}_1^* \mathbf{U}_1^{*\top}$. Then by Equation (16), $\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top = \mathbf{U}_t^* \mathbf{U}_t^{*\top}$ for all $t \in [T]$. Lastly, since $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, we have that $\nabla_{\mathbf{A}} \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, so

$$\hat{\mathbf{A}} = \mathbf{A}^* + \frac{1}{T} \sum_{t=1}^T \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{U}_t \mathbf{U}_t^\top = \mathbf{A}^*$$

□

B.5 Proof of Theorem 4

Proof. Clearly if $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$, then $(\hat{\mathbf{A}}, \hat{\mathbf{U}})$ is an SOSP. The reverse direction is the challenging part of the proof. We equivalently prove that if $(\hat{\mathbf{A}}, \hat{\mathbf{U}})$ is a critical point and $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) \neq 0$, then $\nabla^2 \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}})$ has a negative eigenvalue.

Assume for the sake of contradiction that $(\hat{\mathbf{A}}, \hat{\mathbf{U}})$ is a critical point and $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) \neq 0$. Then,

$$\nabla_{\mathbf{A}} \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = T(\hat{\mathbf{A}} - \mathbf{A}^*) + \sum_{t=1}^T (\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top}) = \mathbf{0} \quad (30)$$

$$\nabla_{\mathbf{U}_t} \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 2 \left(\hat{\mathbf{A}} - \mathbf{A}^* + \hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top} \right) \hat{\mathbf{U}}_t = \mathbf{0} \quad (31)$$

Thus,

$$\hat{\mathbf{A}} = \mathbf{A}^* - \frac{1}{T} \sum_{t=1}^T (\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top}). \quad (32)$$

Define $\mathbf{B}_t(\hat{\mathbf{U}}) = \hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \frac{1}{T} \sum_{s=1}^T (\hat{\mathbf{U}}_s \hat{\mathbf{U}}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top})$. Despite being a slight abuse of notation, we refer to $\mathbf{B}_t(\hat{\mathbf{U}})$ as just \mathbf{B}_t for the remainder of the proof.

Then (31) equivalently states:

$$\mathbf{B}_t \hat{\mathbf{U}}_t = \mathbf{0}. \quad (33)$$

Note that by construction, $\sum_{t=1}^T \mathbf{B}_t = \mathbf{0}$.

Considering \mathcal{L} as a function of the flattened vector $[\text{vec}(\mathbf{A}); \text{vec}(\mathbf{U}_1); \text{vec}(\mathbf{U}_2)]$, and let $\mathbf{U}_1 = [\mathbf{x}_1 \dots \mathbf{x}_k]$, $\mathbf{U}_2 = [\mathbf{y}_1 \dots \mathbf{y}_k]$, we compute the Hessian

$$\nabla^2 \mathcal{L} = \begin{bmatrix} \nabla_{\mathbf{A}}^2 \mathcal{L} & \nabla_{\mathbf{U}_1} \nabla_{\mathbf{A}} \mathcal{L} & \nabla_{\mathbf{U}_2} \nabla_{\mathbf{A}} \mathcal{L} \\ (\nabla_{\mathbf{U}_1} \nabla_{\mathbf{A}} \mathcal{L})^\top & \nabla_{\mathbf{U}_1}^2 \mathcal{L} & \mathbf{0} \\ (\nabla_{\mathbf{U}_2} \nabla_{\mathbf{A}} \mathcal{L})^\top & \mathbf{0} & \nabla_{\mathbf{U}_2}^2 \mathcal{L} \end{bmatrix} \quad (34)$$

where

$$\begin{aligned}
\nabla_{\hat{\mathbf{A}}}^2 \mathcal{L}^* &= 2\mathbf{I}_{d^2} \\
\nabla_{\mathbf{U}_1} \nabla_{\hat{\mathbf{A}}} \mathcal{L}^* &= [(\mathbf{x}_1 \oplus \mathbf{x}_1) \dots (\mathbf{x}_k \oplus \mathbf{x}_k)] \in \mathbb{R}^{d^2 \times dk} \\
\nabla_{\mathbf{U}_2} \nabla_{\hat{\mathbf{A}}} \mathcal{L}^* &= [(\mathbf{y}_1 \oplus \mathbf{y}_1) \dots (\mathbf{y}_k \oplus \mathbf{y}_k)] \in \mathbb{R}^{d^2 \times dk} \\
\nabla_{\mathbf{U}_1}^2 \mathcal{L}^* &= 2(\mathbf{A} + \mathbf{U}_1 \mathbf{U}_1^\top - \mathbf{A}^* - \mathbf{U}_1^* \mathbf{U}_1^{*\top}) \otimes \mathbf{I}_k \\
&\quad + 2 \begin{bmatrix} \mathbf{x}_1 \mathbf{x}_1^\top + \|\mathbf{x}_1\|_2^2 \mathbf{I} & \mathbf{x}_1^\top \mathbf{x}_2 \mathbf{I} + \mathbf{x}_2 \mathbf{x}_1^\top & \dots & \mathbf{x}_1^\top \mathbf{x}_k \mathbf{I} + \mathbf{x}_k \mathbf{x}_1^\top \\ \mathbf{x}_2^\top \mathbf{x}_1 \mathbf{I} + \mathbf{x}_1 \mathbf{x}_2^\top & \mathbf{x}_2 \mathbf{x}_2^\top + \|\mathbf{x}_2\|_2^2 \mathbf{I} & \dots & \mathbf{x}_2^\top \mathbf{x}_k \mathbf{I} + \mathbf{x}_k \mathbf{x}_2^\top \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{x}_k^\top \mathbf{x}_1 \mathbf{I} + \mathbf{x}_1 \mathbf{x}_k^\top & \dots & \dots & \mathbf{x}_k \mathbf{x}_k^\top + \|\mathbf{x}_k\|_2^2 \mathbf{I} \end{bmatrix} \\
\nabla_{\mathbf{U}_2}^2 \mathcal{L}^* &= 2(\mathbf{A} + \mathbf{U}_2 \mathbf{U}_2^\top - \mathbf{A}^* - \mathbf{U}_2^* \mathbf{U}_2^{*\top}) \otimes \mathbf{I}_k \\
&\quad + 2 \begin{bmatrix} \mathbf{y}_1 \mathbf{y}_1^\top + \|\mathbf{y}_1\|_2^2 \mathbf{I} & \mathbf{y}_1^\top \mathbf{y}_2 \mathbf{I} + \mathbf{y}_2 \mathbf{y}_1^\top & \dots & \mathbf{y}_1^\top \mathbf{y}_k \mathbf{I} + \mathbf{y}_k \mathbf{y}_1^\top \\ \mathbf{y}_2^\top \mathbf{y}_1 \mathbf{I} + \mathbf{y}_1 \mathbf{y}_2^\top & \mathbf{y}_2 \mathbf{y}_2^\top + \|\mathbf{y}_2\|_2^2 \mathbf{I} & \dots & \mathbf{y}_2^\top \mathbf{y}_k \mathbf{I} + \mathbf{y}_k \mathbf{y}_2^\top \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}_k^\top \mathbf{y}_1 \mathbf{I} + \mathbf{y}_1 \mathbf{y}_k^\top & \dots & \dots & \mathbf{y}_k \mathbf{y}_k^\top + \|\mathbf{y}_k\|_2^2 \mathbf{I} \end{bmatrix}
\end{aligned}$$

Note that \oplus denotes the Kronecker sum defined as $\mathbf{X} \oplus \mathbf{Y} = \mathbf{I} \otimes \mathbf{X} + \mathbf{Y} \otimes \mathbf{I}$ where \otimes is the Kronecker product.

Lemma 3. $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ if and only if $\mathbf{B}_t = \mathbf{0}$ for each $t \in [T]$.

Proof. Since $(\hat{\mathbf{A}}, \hat{\mathbf{U}})$ is a critical point, then plugging Equation (32) into the definition of \mathcal{L} gives that

$$\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = \frac{1}{2} \sum_{t=1}^T \|\mathbf{B}_t\|_F^2.$$

Thus $\mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) = 0$ if and only if $\mathbf{B}_t = \mathbf{0} \quad \forall t$. □

Lemma 4. If $\nabla_{\hat{\mathbf{U}}}^2 \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) \succcurlyeq \mathbf{0}$, then the eigenvectors corresponding to the non-zero eigenvalues of $\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top$ are the leading non-negative eigenvectors of $\mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \hat{\mathbf{A}}$ for all $t \in [T]$.

Proof. Consider the function $\bar{f}_t(\mathbf{U}_t; \hat{\mathbf{A}}) = \frac{1}{2} \|\mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \hat{\mathbf{A}} - \mathbf{U}_t \mathbf{U}_t^\top\|_F^2$. \bar{f}_t is simply the t th summand in \mathcal{L}^* where $\mathbf{A} = \hat{\mathbf{A}}$ is fixed and we only consider the variable \mathbf{U}_t . Minimizing \bar{f}_t is identical to the problem of symmetric matrix factorization.

Using well-known properties of symmetric matrix factorization, since $\nabla \bar{f}_t(\hat{\mathbf{U}}_t) = \mathbf{0}$, we must have that $\hat{\mathbf{U}}_t = \mathbf{V}_t \mathbf{\Gamma}$ where the columns of \mathbf{V}_t are the properly scaled eigenvectors of $\mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \hat{\mathbf{A}}$ with non-negative eigenvalues where each column has norm equal to the square root of its corresponding eigenvalue, and $\mathbf{\Gamma} \in O_k$ is some orthogonal matrix. Further, if the eigenvectors corresponding to the non-zero eigenvalues of $\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top$ are not the leading non-negative eigenvectors, then $\nabla^2 \bar{f}_t(\hat{\mathbf{U}}) \not\preccurlyeq \mathbf{0}$ by [36]. Since $\nabla^2 \bar{f}_t(\hat{\mathbf{U}}_t)$ is a diagonal block of $\nabla^2 \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}})$, $\nabla^2 \bar{f}_t(\hat{\mathbf{U}}_t) \not\preccurlyeq \mathbf{0}$ would imply $\nabla^2 \mathcal{L}^*(\hat{\mathbf{A}}, \hat{\mathbf{U}}) \not\preccurlyeq \mathbf{0}$. □

Remark 3. Without loss of generality, we can assume that the eigenvectors corresponding to the non-zero eigenvalues of $\hat{\mathbf{U}}_t \hat{\mathbf{U}}_t^\top$ are the leading non-negative eigenvectors of $\mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \hat{\mathbf{A}}$ for all i .

Lemma 5. $(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top) \mathbf{x} = (\mathbf{U}_2^* \mathbf{U}_2^{*\top} - \mathbf{U}_1^* \mathbf{U}_1^{*\top}) \mathbf{x}$ for all $\mathbf{x} \in \text{im}(\hat{\mathbf{U}}_1) + \text{im}(\hat{\mathbf{U}}_2)$.

Proof. Recall $B_1 = \frac{1}{2} \left(\hat{U}_1 \hat{U}_1^\top - U_1^* U_1^{*\top} - \hat{U}_2 \hat{U}_2^\top + U_2^* U_2^{*\top} \right)$. Then applying first-order stationarity and the fact that $B_2 = -B_1$, we have

$$\begin{aligned} (\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top) \hat{U}_1 &= (U_2^* U_2^{*\top} - U_1^* U_1^{*\top}) \hat{U}_1 \\ (\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top) \hat{U}_2 &= (U_2^* U_2^{*\top} - U_1^* U_1^{*\top}) \hat{U}_2. \end{aligned}$$

□

Corollary 5. $\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top$ and $U_2^* U_2^{*\top} - U_1^* U_1^{*\top}$ share an eigenbasis.

Proof. Using the lemma, any non-zero eigenvector-eigenvalue pair of $\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top$ is also an eigenvector-eigenvalue pair of $U_2^* U_2^{*\top} - U_1^* U_1^{*\top}$. Denote the space defined by the span of these eigenvectors as \mathcal{S} . Then all other eigenvectors of $U_2^* U_2^{*\top} - U_1^* U_1^{*\top}$ are orthogonal to \mathcal{S} , so they are also 0-eigenvectors of $\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top$. Thus the two matrices share an eigenbasis. □

Corollary 6. $\dim(\text{im } \hat{U}_1 + \text{im } \hat{U}_2) \leq 2k - 1$, i.e., the set of columns of \hat{U}_1 and \hat{U}_2 are not linearly independent.

Proof. Assume for contradiction that the vectors in the set $\mathcal{S} = \{\hat{U}_1 e_i \mid i = 1, \dots, k\} \cup \{\hat{U}_2 e_i \mid i = 1, \dots, k\}$ are linearly independent, where e_i is the i th standard basis vector in \mathbb{R}^k .

Then note that $(\hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top) x \neq \mathbf{0}$ and $(U_1^* U_1^{*\top} - U_2^* U_2^{*\top}) x \neq \mathbf{0}$ for all $x \in \mathcal{S}$. By Lemma (5), $\hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top$ and $U_1^* U_1^{*\top} - U_2^* U_2^{*\top}$ agree for each vector on the $2k$ -dimensional space $\text{span}(\mathcal{S})$. But, both $\text{rank}(\hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top), \text{rank}(U_1^* U_1^{*\top} - U_2^* U_2^{*\top}) \leq 2k$ by construction. Then by dimension counting, $\hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top$ and $U_1^* U_1^{*\top} - U_2^* U_2^{*\top}$ must send $\text{span}\{\mathcal{S}\}^\perp$ to $\mathbf{0}$. Thus, $\hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top$ and $U_1^* U_1^{*\top} - U_2^* U_2^{*\top}$ agree on the entire basis formed by concatenating basis vectors of $\text{span}\{\mathcal{S}\}^\perp$ with those of $\text{span}(\mathcal{S})$. This implies that $\hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top = U_1^* U_1^{*\top} - U_2^* U_2^{*\top}$ and thus $B_1 = \hat{U}_1 \hat{U}_1^\top - \hat{U}_2 \hat{U}_2^\top - U_1^* U_1^{*\top} + U_2^* U_2^{*\top} = \mathbf{0}$. Then $B_2 = -B_1 = \mathbf{0}$ so by Lemma 3, $\mathcal{L}^*(\hat{A}, \hat{U}) = 0$ which is a contradiction. □

Lemma 6. $U_2^* U_2^{*\top} - U_1^* U_1^{*\top}$ has exactly k positive and k negative eigenvalues.

Proof. First, note that $U_2^* U_2^{*\top}$ has exactly k positive eigenvalues and $k - d$ eigenvalues of $\mathbf{0}$. Then $U_2^* U_2^{*\top} - (U_1^* e_1)(U_1^* e_1)^\top$ has rank $k + 1$ because of the linear independence of the columns of the combined set of columns U_1^* and U_2^* . Further, since we subtract $(U_1^* e_1)(U_1^* e_1)^\top$, we must be accumulating an additional negative eigenvalue relative to $U_2^* U_2^{*\top}$. Continuing this process shows that subtracting $(U_1^* e_{j+1})(U_1^* e_{j+1})^\top$ from $U_2^* U_2^{*\top} - \sum_{t=1}^j (U_1^* e_t)(U_1^* e_t)^\top$ contributes exactly one more negative eigenvalue, since $U_1^* e_{j+1}$ can never be written as a linear combination of $\{U_1^* e_1, \dots, U_1^* e_k, U_2^* e_1, \dots, U_2^* e_j\}$ for $0 < j < k$. The result then follows from induction. □

Lemma 7. $\text{rank}(\hat{U}_1) = \text{rank}(\hat{U}_2) = k$.

Proof. Assume for contradiction that $\text{rank}(\hat{U}_1) = m < k$ without loss of generality. Since by Remark (3) we assume the columns of \hat{U}_1 are the leading k non-negative eigenvectors of $A^* + U_1^* U_1^{*\top} - \hat{A} = \hat{U}_1 \hat{U}_1^\top - B_1$, this must imply that $A^* + U_1^* U_1^{*\top} - \hat{A} - \hat{U}_1 \hat{U}_1^\top = -B_1 \preceq \mathbf{0}$.

Plugging in the definition of B_1 gives that $\frac{1}{2} \left(\hat{U}_1 \hat{U}_1^\top - U_1^* U_1^{*\top} - \hat{U}_2 \hat{U}_2^\top + U_2^* U_2^{*\top} \right) \succeq \mathbf{0}$. Thus, $\hat{U}_1 \hat{U}_1^\top \succeq U_1^* U_1^{*\top} + \hat{U}_2 \hat{U}_2^\top - U_2^* U_2^{*\top} \succeq U_1^* U_1^{*\top} - U_2^* U_2^{*\top}$. This contradicts the fact from Lemma (6) that $U_1^* U_1^{*\top} - U_2^* U_2^{*\top}$ has k positive eigenvalues. □

With this lemma, we will prove the existence of a direction of $\nabla^2 \mathcal{L}^*$ with negative curvature. Instead of directly working with this matrix, we instead use the Schur complement to work with a different form.

Theorem 5. (Schur Complement) Since $\nabla_{\hat{A}}^2 \mathcal{L}^*(\hat{A}, \hat{U}) = 2\mathbf{I} \succ \mathbf{0}$, $\nabla^2 \mathcal{L}^*(\hat{A}, \hat{U}) \succcurlyeq \mathbf{0}$ if and only if $\nabla_{\hat{U}}^2 \mathcal{L}^*(\hat{A}, \hat{U}) - \left(\nabla_{\hat{A}} \nabla_{\hat{U}} \mathcal{L}^*(\hat{A}, \hat{U}) \right) \left(\nabla_{\hat{A}}^2 \mathcal{L}^*(\hat{A}, \hat{U}) \right)^{-1} \left(\nabla_{\hat{U}} \nabla_{\hat{A}} \mathcal{L}^*(\hat{A}, \hat{U}) \right) \succcurlyeq \mathbf{0}$.

Define $M = \nabla_{\hat{U}}^2 \mathcal{L}^*(\hat{A}, \hat{U}) - \left(\nabla_{\hat{A}} \nabla_{\hat{U}} \mathcal{L}^*(\hat{A}, \hat{U}) \right) \left(\nabla_{\hat{A}}^2 \mathcal{L}^*(\hat{A}, \hat{U}) \right)^{-1} \left(\nabla_{\hat{U}} \nabla_{\hat{A}} \mathcal{L}^*(\hat{A}, \hat{U}) \right)$.

For example, when $k = 2$ and letting $\mathbf{U}_1 = [\mathbf{x}_1 \ \mathbf{x}_2]$, $\mathbf{U}_2 = [\mathbf{y}_1 \ \mathbf{y}_2]$, we have

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{12}^\top & M_{22} \end{bmatrix},$$

where

$$M_{11} = \begin{bmatrix} 2\mathbf{B}_1 + \mathbf{x}_1 \mathbf{x}_1^\top + \|\mathbf{x}_1\|_2^2 & \mathbf{x}_1^\top \mathbf{x}_2 \mathbf{I} + \mathbf{x}_2 \mathbf{x}_1^\top \\ \mathbf{x}_2^\top \mathbf{x}_1 \mathbf{I} + \mathbf{x}_1 \mathbf{x}_2^\top & 2\mathbf{B}_1 + \mathbf{x}_2 \mathbf{x}_2^\top + \|\mathbf{x}_2\|_2^2 \end{bmatrix}$$

$$M_{12} = \begin{bmatrix} -\mathbf{x}_1^\top \mathbf{y}_1 \mathbf{I} - \mathbf{y}_1 \mathbf{x}_1^\top & -\mathbf{x}_1^\top \mathbf{y}_2 \mathbf{I} - \mathbf{y}_2 \mathbf{x}_1^\top \\ -\mathbf{x}_2^\top \mathbf{y}_1 \mathbf{I} - \mathbf{y}_1 \mathbf{x}_2^\top & \mathbf{x}_2^\top \mathbf{y}_2 \mathbf{I} - \mathbf{y}_2 \mathbf{x}_2^\top \end{bmatrix}$$

$$M_{22} = \begin{bmatrix} 2\mathbf{B}_2 + \mathbf{y}_1 \mathbf{y}_1^\top + \|\mathbf{y}_1\|_2^2 & \mathbf{y}_1^\top \mathbf{y}_2 \mathbf{I} + \mathbf{y}_2 \mathbf{y}_1^\top \\ \mathbf{y}_2^\top \mathbf{y}_1 \mathbf{I} + \mathbf{y}_1 \mathbf{y}_2^\top & 2\mathbf{B}_2 + \mathbf{y}_2 \mathbf{y}_2^\top + \|\mathbf{y}_2\|_2^2 \end{bmatrix}$$

For brevity, we do not include the full form of M for general k . However, we can make an easy simplification that will allow for a much cleaner expression.

Using Corollaries (5) and (6), there is an eigenvector \mathbf{z} of $\mathbf{U}_2^* \mathbf{U}_2^{*\top} - \mathbf{U}_1^* \mathbf{U}_1^{*\top}$ with eigenvalue $\lambda \neq 0$ such that $\mathbf{z} \in \ker \left(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top \right)$. Assume without loss of generality that $\lambda > 0$, and consider $\boldsymbol{\alpha} \in \mathbb{R}^{2k}$. Define the function $g(\cdot; \mathbf{z}) : \mathbb{R}^{2k} \rightarrow \mathbb{R}$ parameterized by \mathbf{z} such that $g(\boldsymbol{\alpha}; \mathbf{z}) = (\boldsymbol{\alpha} \otimes \mathbf{z})^\top M (\boldsymbol{\alpha} \otimes \mathbf{z})$, where we partition $\boldsymbol{\alpha} = [\boldsymbol{\alpha}_1; \boldsymbol{\alpha}_2]$, $\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2 \in \mathbb{R}^k$. Then after some algebra,

$$g(\boldsymbol{\alpha}; \mathbf{z}) = \left\| \hat{\mathbf{U}}_1 \boldsymbol{\alpha}_1 + \hat{\mathbf{U}}_2 \boldsymbol{\alpha}_2 \right\|_2^2 + \lambda \left(\|\boldsymbol{\alpha}_1\|_2^2 - \|\boldsymbol{\alpha}_2\|_2^2 \right). \quad (35)$$

We prove the existence of $\boldsymbol{\alpha} \in \mathbb{R}^{2k}$, $\mathbf{x} \in \mathbb{R}^d$ such that $g(\boldsymbol{\alpha}; \mathbf{x}) < 0$ considering two different cases. Define $N^- : S_d \rightarrow \mathbb{Z}$ as the function that returns the number of strictly negative eigenvalues of its input.

Case 1: $N^- \left(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top \right) < k$.

Using Corollary (6), we can pick $\boldsymbol{\alpha}$ such that $\hat{\mathbf{U}}_1 \boldsymbol{\alpha}_1 + \hat{\mathbf{U}}_2 \boldsymbol{\alpha}_2 = \mathbf{0}$, $\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2 \neq \mathbf{0}$.

Because $N^- \left(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top \right) < k$, $N^- \left(\mathbf{U}_2^* \mathbf{U}_2^{*\top} - \mathbf{U}_1^* \mathbf{U}_1^{*\top} \right) = k$, and $\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top$ and $\mathbf{U}_2^* \mathbf{U}_2^{*\top} - \mathbf{U}_1^* \mathbf{U}_1^{*\top}$ share an eigenbasis by Corollary 5, there exists $\mathbf{z}^- \in \mathbb{R}^d$ that is a λ^- -eigenvector of $\mathbf{U}_2^* \mathbf{U}_2^{*T} - \mathbf{U}_1^* \mathbf{U}_1^{*T}$, $\lambda^- < 0$, where $\mathbf{z} \in \ker \left(\hat{\mathbf{U}}_2 \hat{\mathbf{U}}_2^\top - \hat{\mathbf{U}}_1 \hat{\mathbf{U}}_1^\top \right)$

Then for the same choice of $\boldsymbol{\alpha}$,

$$\begin{aligned} \text{sign}(g(\boldsymbol{\alpha}; \mathbf{z})) &= \text{sign} \left(\|\boldsymbol{\alpha}_1\|_2^2 - \|\boldsymbol{\alpha}_2\|_2^2 \right) \\ \text{sign}(g(\boldsymbol{\alpha}; \mathbf{z}^-)) &= \text{sign} \left(\|\boldsymbol{\alpha}_2\|_2^2 - \|\boldsymbol{\alpha}_1\|_2^2 \right). \end{aligned}$$

Then if $\|\boldsymbol{\alpha}_1\|_2 \neq \|\boldsymbol{\alpha}_2\|_2$, one of the above expressions is negative and thus M has a negative eigenvalue. This then implies $\nabla^2 \mathcal{L}^*(\hat{A}, \hat{U}) \not\succeq \mathbf{0}$.

Otherwise $\|\alpha_1\|_2 = \|\alpha_2\|_2$. Then $g(\alpha; \mathbf{z}) = 0$, but $\nabla_{\alpha_1} g(\alpha; \mathbf{z}) = \hat{U}_1^\top \left(\hat{U}_1 \bar{\alpha}_1 + \hat{U}_2 \alpha_2 \right) - 2\lambda \alpha_2 = -2\lambda \alpha_2 \neq 0$. Thus $g(\alpha; \mathbf{z}) = 0$ and $\nabla g(\alpha; \mathbf{z}) \neq 0$ so there exists $\bar{\alpha}$ in an infinitesimal neighborhood around α where $g(\bar{\alpha}; \mathbf{z}) < 0$. Thus M has a negative eigenvalue so $\nabla^2 \mathcal{L}^*(\hat{A}, \hat{U}) \not\preceq 0$.

Case 2: $N^-(\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top) = k$.

Define $m = \dim \left(\text{im}(\hat{U}_1) \cap \text{im}(\hat{U}_2) \right)$. By Corollary 6, $m \geq 1$, so we can select orthogonal matrix $\Gamma \in O_k$ such that $\hat{U}_2 \Gamma e_1 \in \left(\text{im}(\hat{U}_1) \cap \text{im}(\hat{U}_2) \right)$. Define $\mathbf{y} = \hat{U}_2 \Gamma e_1$.

Clearly for any $\mathbf{B} \in S_d$ and $\mathbf{R} \in S_d^+$, $N^-(\mathbf{B}) \geq N^-(\mathbf{B} + \mathbf{R})$. Then since $N^-(-\hat{U}_1 \hat{U}_1^\top) = k$ by Lemma (7), we have that

$$\begin{aligned} k &= N^-(-\hat{U}_1 \hat{U}_1^\top) \geq N^-(\mathbf{y} \mathbf{y}^\top - \hat{U}_1 \hat{U}_1^\top) = N^-(\left(\hat{U}_2 \Gamma e_1 \right) \left(\hat{U}_2 \Gamma e_1 \right)^\top - \hat{U}_1 \hat{U}_1^\top) \\ &\geq N^-(\left(\hat{U}_2 \Gamma \right) \left(\hat{U}_2 \Gamma \right)^\top - \hat{U}_1 \hat{U}_1^\top) = N^-(\hat{U}_2 \hat{U}_2^\top - \hat{U}_1 \hat{U}_1^\top) = k, \end{aligned}$$

Thus, $N^-(\mathbf{y} \mathbf{y}^\top - \hat{U}_1 \hat{U}_1^\top) = k$. But, since $\mathbf{y} \in \text{im}(\hat{U}_1)$, $\text{rank} \left(\mathbf{y} \mathbf{y}^\top - \hat{U}_1 \hat{U}_1^\top \right) = k$. Thus,

$$\mathbf{y} \mathbf{y}^\top - \hat{U}_1 \hat{U}_1^\top \preceq 0. \quad (36)$$

Take α such that $\hat{U}_1 \alpha_1 = -\mathbf{y}$ and $\alpha_2 = \Gamma e_1$. Then

$$\mathbf{y}_1 \mathbf{y}_1^\top - \hat{U}_1 \hat{U}_1^\top = \left(\hat{U}_1 \alpha \right) \left(\hat{U}_1 \alpha \right)^\top - \hat{U}_1 \hat{U}_1^\top \quad (37)$$

$$= \hat{U}_1 \left(\alpha_1 \alpha_1^\top - \mathbf{I} \right) \hat{U}_1^\top \preceq 0. \quad (38)$$

Therefore $\|\alpha_1\|_2 \leq 1$.

Then $g(\alpha; \mathbf{z}) = \left\| \hat{U}_1 \alpha_1 + \hat{U}_2 \alpha_2 \right\|_2^2 + \lambda \left(\|\alpha_1\|_2^2 - \|\alpha_2\|_2^2 \right) = \lambda \left(\|\alpha_1\|_2^2 - 1 \right) \leq 0$.

If $g(\alpha; \mathbf{z}) < 0$, then we are done. Otherwise, $g(\alpha; \mathbf{z}) = 0$. Then the same analysis from Case 1 will show that $\nabla g(\alpha; \mathbf{z}) \neq 0$, so there exists $\bar{\alpha}$ in an infinitesimal neighborhood around α where $g(\bar{\alpha}; \mathbf{z})$ is strictly negative. This then implies our desired result. \square

B.6 Derivation of Equation (8)

Recall our generative model where each input sample $\mathbf{x} \in \mathbb{R}^d$ satisfies $\mathbb{E}[\mathbf{x}] = \mathbf{0}$ and $\mathbb{E}[\mathbf{x} \mathbf{x}^\top] = \sigma_x^2 \mathbf{I}_d$, each noise sample is generated independently of \mathbf{x} as $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I}_d)$, and $\mathbf{y} = (\mathbf{A}^* + \mathbf{R}_t^*) \mathbf{x} + \epsilon$. Then,

$$\begin{aligned} 2\mathbb{E}[\mathcal{L}_t^1(\mathbf{A})] &= \mathbb{E} \left[\|\mathbf{y} - \mathbf{A} \mathbf{x}\|_2^2 \right] \\ &= \mathbb{E} \left[\|(\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}) \mathbf{x} + \epsilon\|_2^2 \right] \\ &= \mathbb{E} \left[\|(\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \mathbf{x}\|_2^2 + \|\epsilon\|_2^2 + 2\epsilon^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \mathbf{x} \right] \\ &= \mathbb{E} \left[\text{tr} \left(\mathbf{x}^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t)^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \mathbf{x} \right) \right] + \sigma_\epsilon^2 \\ &\quad + 2\mathbb{E}[\epsilon]^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \mathbb{E}[\mathbf{x}] \\ &= \mathbb{E} \left[\text{tr} \left\{ (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t)^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \mathbf{x} \mathbf{x}^\top \right\} \right] + \sigma_\epsilon^2 \\ &= \text{tr} \left\{ (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t)^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \mathbb{E}[\mathbf{x} \mathbf{x}^\top] \right\} + \sigma_\epsilon^2 \\ &= \sigma_x^2 \text{tr} \left((\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t)^\top (\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t) \right) + \sigma_\epsilon^2 \\ &= \sigma_x^2 \|\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t\|_F^2 + \sigma_\epsilon^2 \end{aligned}$$

Thus, $\mathbb{E}[\mathcal{L}_t^1(\mathbf{A}_t)] = \frac{1}{2} \left(\sigma_x^2 \|\mathbf{A}^* + \mathbf{R}_t^* - \mathbf{A}_t\|_F^2 + \sigma_\epsilon^2 \right)$. Then $\mathbb{E}[\mathcal{L}_t^{n_t}(\mathbf{A}_t)] = \mathbb{E}[\mathcal{L}_t^1(\mathbf{A}_t)]$ by linearity of expectation, so

$$\frac{1}{\sigma_x^2} \left(\mathbb{E}[\mathcal{L}_t^{n_t}(\mathbf{A}_t)] - \frac{\sigma_\epsilon^2}{2} \right) = \frac{1}{2} \left\| \mathbf{A}^* + \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \mathbf{A}_t \right\|_F^2$$

C Limitations of Standard Retraining for Asymmetric Adapters

In this section we show that even if the ground truth rank- k adapters \mathbf{R}_t^* are asymmetric in our theoretical model defined in Section 3, standard retraining still fails to recover a retraining solution $\hat{\mathbf{A}}_{\text{SR}}$ that guarantees that either $\text{rank}(\hat{\mathbf{A}}_{\text{SR}} - \mathbf{A}^*)$ or $\|\hat{\mathbf{A}}_{\text{SR}} - \mathbf{A}^*\|$ is small.

Consider task-specific adapters $\mathbf{R}_t^* \in \mathbb{R}^{d \times d}$ drawn independently from a general distribution \mathcal{D}_R such that $\text{rank}(\mathbf{R}_t^*) = k$ and \mathcal{D}_R is absolutely continuous with respect to the Lebesgue measure. Then since $\hat{\mathbf{A}}_{\text{SR}} = \mathbf{A}^* + \frac{1}{T} \sum_{t=1}^T \mathbf{R}_t^*$, we immediately have that $\hat{\mathbf{A}}_{\text{SR}}$ and \mathbf{A}^* are far apart in terms of rank:

$$\text{rank}(\hat{\mathbf{A}}_{\text{SR}} - \mathbf{A}^*) = \min\{d, kT\} \quad \text{w.p. } 1.$$

As a result, the necessary adaptation rank during fine-tuning to realize the test task parameters $\mathbf{A}^* + \mathbf{R}_{T+1}^*$ is $\min\{d, k(T+1)\}$, just as in Proposition 1.

Additionally, we show that we can never guarantee that $\|\hat{\mathbf{A}}_{\text{SR}} - \mathbf{A}^*\|$ is small for general adapter distributions \mathcal{D}_R . Let mean of the adapter distribution be $M_R = \mathbb{E}[\mathbf{R}_t^*]$. Since the error between $\hat{\mathbf{A}}_{\text{SR}}$ and \mathbf{A}^* is the average of retraining adapters $\frac{1}{T} \sum_{t=1}^T \mathbf{R}_t^*$, then by the law of large numbers we must have that as the number of retraining tasks $T \rightarrow \infty$, then $\hat{\mathbf{A}}_{\text{SR}} - \mathbf{A}^* \rightarrow M_R$. Thus, we can only guarantee that standard retraining will approach a reasonable solution in the special and impractical case where both (i) the number of tasks approaches infinity, and (ii) the mean of the adapter distribution $M_R = \mathbf{0}$. If either of these conditions do not hold, then standard retraining will provably fail to recover \mathbf{A}^* .

We lastly note that even though we model each task t by the additive parameterization $\mathbf{A}^* + \mathbf{R}_t^*$, we cannot absorb the mean M_R into \mathbf{A}^* to reduce to the case where $M_R = \mathbf{0}$ without destroying the low-rank structure of \mathbf{R}_t^* . Specifically, given a problem instance defined by tasks \mathcal{T}_t parameterized by $\mathbf{A}^* + \mathbf{R}_t^*$ where $M_R = \mathbb{E}[\mathbf{R}_t^*]$, we could equivalently consider the tasks as parameterized by $(\mathbf{A}^* + M_R) + (\mathbf{R}_t^* - M_R)$. Then the transformed adapters $\mathbf{R}_t^* - M_R$ are zero-mean, but $\text{rank}(\mathbf{R}_t^* - M_R)$ may be as large as the ambient dimension d , eliminating the low-rank structure needed for the computational and sample-efficiency gains of low-rank adaptation during test-time fine-tuning.

D Synthetic Experiments

D.1 Linear Model

We first test our meta-learning framework on synthetic regression tasks. We consider data $\mathbf{y}_{t,j} = (\mathbf{A}^* + \mathbf{R}_t^*) \mathbf{x}_{t,j} + \epsilon_{t,j}$ for task t and $j \in [n_t]$, where we generate $\mathbf{x}_{t,j}$, $\epsilon_{t,j}$, and \mathbf{R}_t^* as in Section 3 with $\mathbf{x}_{t,j} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\epsilon_{t,j} \sim \mathcal{N}(0, .01 \times \mathbf{I})$, and $\mathbf{R}_t^* = \mathbf{U}_t^* \mathbf{U}_t^{*\top}$ with each element of \mathbf{U}_t^* sampled as a $\mathcal{N}(0, 1)$ random variable. We similarly generate the entries of \mathbf{A}^* as $\mathcal{N}(0, 1)$ random variables. We set the number of samples per retraining task for each task t as $n_t = N$, so each retraining task is equipped with N samples. Further, we denote the number of test task samples $n_{T+1} = n$. We set \mathcal{L} to be the mean squared error loss, and we run gradient descent on the standard retraining (1) and the LoRA-ML (5) objectives. For the LoRA-ML objective, we alternate between the outer step, updating the shared parameter \mathbf{A} , and the inner steps, where we update the task specific parameter \mathbf{U}_t for each task independently. We run 3000 outer steps, where after each outer step we run 10 inner steps. We use a learning rate of 3×10^{-2} for the outer steps and 3×10^{-3} for the inner steps. For the standard retraining objective, we simply run gradient descent using the learning rate 3×10^{-2} over 3000 epochs.

After recovering \hat{A} during retraining, we apply the low-rank adaptation QV^\top when fine-tuning to the test task. When $T = 2$, we use a rank- $3k$ adaptation during fine-tuning to account for the inexact recovery explained in Theorem 2, and otherwise use a rank- k adaptation. To perform the adaptation, we run gradient descent on the LoRA objective (9) using a learning rate of 5×10^{-3} .

In each experiment we vary one hyperparameter from a fixed set of values and plot the prediction error between the recovered model and the ground truth model $\frac{1}{n} \sum_j \|y_{T+1,j} - (\hat{A} + QV^\top)x_{T+1,j}\|_2^2$. Results are averaged over 5 trials, with the shaded region showing the full range of values across trials. For each experiment, we fine-tuned for a different number of epochs, as some required more epochs to converge. The performance by epoch is displayed on the x-axis of each plot.




Figure 2: Linear model fine-tuning performance varying the number of retraining tasks T . This is an enlargement of the left subfigure of Figure 1.




Figure 3: Linear model fine-tuning performance varying the number of samples for the test task n . This is an enlargement of the right subfigure of Figure 1.




Figure 4: Linear model fine-tuning performance varying the number of samples per retraining task N .




Figure 5: Linear model fine-tuning performance varying the ground truth adaptation rank k .




Figure 6: Linear model fine-tuning performance varying the ambient dimension d .

D.2 Shallow Network

We additionally run a synthetic data experiment using data generated using a shallow network architecture. We first generate inputs $\mathbf{x}_{t,j} \in \mathbb{R}^d$ distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. We then generate ground truth parameters $\mathbf{A}^* \in \mathbb{R}^{d \times d}$, $\mathbf{c}^* \in \mathbb{R}^d$, and $\mathbf{U}_t^* \in \mathbb{R}^{d \times k}$ for each task $t \in [T]$ by sampling the entries as $\mathcal{N}(0, 1)$ random variables. We denote $\mathbf{R}_t^* = \mathbf{U}_t^* \mathbf{U}_t^{*\top}$. Lastly, the outputs $y_{t,j} \in \mathbb{R}$ for task t are generated as $y_{t,j} = \mathbf{c}^{*\top} s((\mathbf{A}^* + \mathbf{R}_t^*) \mathbf{x}_{t,j}) + \epsilon_{t,j}$, where $s(\cdot)$ is the element-wise sigmoid function and $\epsilon_{t,j} \sim \mathcal{N}(0, .01)$ is noise. We again define parameters N, n such that $n_t = N$ for all $t \leq T$ and $n_{T+1} = n$. Setting \mathcal{L} to be the mean squared error loss, we apply the AdamW optimizer with a learning rate of 1×10^{-3} to both the standard retraining (1) and the LoRA-ML (5) objectives. After recovering $\hat{\mathbf{A}}$ and $\hat{\mathbf{c}}$ during retraining, we apply the low-rank adaptation \mathbf{QV}^\top only to $\hat{\mathbf{A}}$ when fine-tuning to the test task. In each experiment we vary a single hyperparameter from a fixed set of values and plot the prediction error between the recovered model and the ground truth model $\frac{1}{n} \sum_{j=1}^n \left\| \hat{\mathbf{c}}^\top s((\hat{\mathbf{A}} + \mathbf{QV}^\top) \mathbf{x}_{T+1,j}) - \mathbf{c}^{*\top} s((\mathbf{A}^* + \mathbf{R}_{T+1}^*) \mathbf{x}_{T+1,j}) \right\|_2^2$. To mitigate the effects of outliers, we plot the median over 10 trials, where the tables that accompany each figure report the range of the central 6 values of the last epoch prediction error, clipping the two best and worst trials.

Figures 7, 8, 9, 10, and 11 again show that retraining using the LoRA-ML objective leads to much better fine-tuning performance relative to standard retraining. Figure 7 shows that the LoRA-ML objective effectively exploits the number of tasks, since fine-tuning performance improves as T increases. Further, we observe in Figure 8 that fine-tuning with any number of samples after standard retraining cannot even recover the performance of fine-tuning with just 100 samples after LoRA-ML retraining. We note that the range of the results varied much more in the shallow network setting relative to the linear model experiments. Although LoRA-ML performed much better than standard retraining when looking at the median result, a few individual trials showed outlier results.

The data parameters for both synthetic experiments are summarized in Table 9. Both were performed using a single NVIDIA A40 GPU.




Figure 7: Shallow network fine-tuning performance while varying number of retraining tasks T .

Table 4: Median and central range of the test task loss after fine-tuning (smaller is better) at the last epoch for the shallow network experiment across 10 trials while varying T .

Method	$T = 2$	$T = 3$	$T = 5$
LoRA-ML	0.047 (0.042, 0.127)	0.018 (0.017, 0.028)	0.005 (0.004, 0.024)
SR	0.119 (0.104, 0.161)	0.092 (0.086, 0.225)	0.076 (0.046, 0.107)




Figure 8: Shallow network fine-tuning performance while varying number of test task samples n .

Table 5: Median and central range of the test task loss after fine-tuning (smaller is better) at the last epoch for the shallow network experiment across 10 trials while varying n .

Method	$n = 100$	$n = 1000$	$n = 10000$
LoRA-ML	0.006 (0.005, 0.034)	0.003 (0.002, 0.061)	0.004 (0.004, 0.017)
SR	0.132 (0.114, 0.171)	0.147 (0.109, 0.221)	0.115 (0.112, 0.165)




Figure 9: Shallow network fine-tuning performance while varying number of samples per retraining task N .

Table 6: Median and central range of the test task loss after fine-tuning (smaller is better) at the last epoch for the shallow network experiment across 10 trials while varying N .

Method	$N = 100$	$N = 1000$	$N = 10000$
LoRA-ML	0.138 (0.100, 0.262)	0.010 (0.006, 0.055)	0.023 (0.009, 0.048)
SR	0.354 (0.319, 0.646)	0.075 (0.073, 0.092)	0.085 (0.074, 0.108)




Figure 10: Shallow network fine-tuning performance while varying ground truth adaptation rank k .

Table 7: Median and central range of the test task loss after fine-tuning (smaller is better) at the last epoch for the shallow network experiment across 10 trials while varying k .

Method	$k = 1$	$k = 3$	$k = 5$
LoRA-ML	0.138 (0.100, 0.262)	0.093 (0.077, 0.117)	0.044 (0.043, 0.067)
SR	0.354 (0.319, 0.646)	0.283 (0.261, 0.359)	0.151 (0.140, 0.299)




Figure 11: Shallow network fine-tuning performance while varying ambient dimension d .

Table 8: Median and central range of the test task loss after fine-tuning (smaller is better) at the last epoch for the shallow network experiment across 10 trials while varying d .

Method	$d = 5$	$d = 10$	$d = 20$
LoRA-ML	0.002 (0.001, 0.003)	0.005 (0.004, 0.010)	0.069 (0.050, 0.121)
SR	0.004 (0.003, 0.021)	0.082 (0.082, 0.125)	0.483 (0.438, 0.690)

Table 9: Synthetic Data Parameters

Experiment	T	n	N	k	d	σ_x	σ_ϵ
Linear, varying T	{2,3,5}	100	5000	1	10	1	.1
Linear, varying n	3	{100,1000,10000}	5000	1	10	1	.1
Linear, varying N	3	100	{100,1000,10000}	1	10	1	.1
Linear, varying k	3	100	1000	{1,2,3}	10	1	.1
Linear, varying d	3	100	5000	1	{5,10,20}	1	.1
Shallow Network, varying T	{2,3,5}	100	1000	1	10	1	.1
Shallow Network, varying n	3	{100,1000,10000}	1000	1	10	1	.1
Shallow Network, varying N	3	100	{100,1000,10000}	1	10	1	.1
Shallow Network, varying k	3	100	1000	{1,3,5}	10	1	.1
Shallow Network, varying d	3	100	1000	1	{5,10,20}	1	.1

E Real Data Experiments

In this section, we describe the real data experiments in complete detail. All of the real data experiments were performed on a single NVIDIA H200 GPU.

E.1 Vision Experiments

We use CIFAR-10 [21], and define $T = 4$ binary classification retraining tasks. Each task requires classification between consecutive CIFAR-10 class labels. Specifically, task 1 classifies between classes 1 and 2, task 2 between classes 3 and 4, etc. The test task is binary classification between classes 9 and 10. We evaluate our PEFT-ML framework using a model based on the MLP-Mixer architecture [35]. We use a depth of 1, and an embedding dimension of 512 for the patches. We compare two different adaptation methods: (i) LoRA with rank 1 adapters, and (ii) fine-tuning only the last layer. For both PEFT strategies, we evaluate three retraining strategies: (a) PEFT-ML for

Table 10: Mean test accuracies and standard errors for PEFT-ML and standard retraining methods using LoRA-based and last-layer-based fine-tuning for adapting to a subset of CIFAR-10 classes.

LoRA Fine-Tuning			Last-Layer Fine-Tuning		
Method	Mean	Std. Err.	Method	Mean	Std. Err.
LoRA-ML	86.09	0.35	Last-Layer-ML	83.90	0.17
SR+LoRA	85.29	0.17	SR+Last-Layer	81.16	0.22
Reptile+LoRA	81.30	0.42	Reptile+Last-Layer	72.55	0.39

each PEFT method (LoRA-ML, Last-Layer-ML), (b) standard retraining (SR), and (c) Reptile, a popular gradient-based meta-learning method. We report mean results along with standard errors across 5 trials in Table 10.

For each retraining method, we take the model from the epoch with the best average performance on the validation samples for the retraining tasks to then be fine-tuned. We also apply basic transformations of the retraining data like random cropping and flipping to reduce overfitting. During fine-tuning, we take the best accuracy on the heldout samples for the test task across epochs for each trial, and we report the average of these values in the Table 10. We summarize the hyperparameter choices in Tables 11 and 12.

Table 11: Retraining hyperparameters for vision experiments.

Hyperparameter	Standard Retraining	LoRA-ML	Last-Layer-ML	Reptile
Learning Rate	10^{-4}	10^{-4}	10^{-4}	10^{-2}
Outer Learning Rate	N/A	N/A	N/A	10^{-5}
Weight Decay	10^{-5}	10^{-5}	10^{-5}	10^{-5}
Epochs	100	100	100	5
Inner Epochs	N/A	N/A	N/A	20
Optimizer	Adam	Adam	Adam	Adam
Batch Size	256	256	256	256
LoRA Rank	N/A	1	N/A	N/A

Table 12: Fine-Tuning hyperparameters for vision experiments

Hyperparameter	LoRA Fine-Tuning	Last-Layer Fine-Tuning
Learning Rate	$5 \cdot 10^{-4}$	$5 \cdot 10^{-4}$
Weight Decay	10^{-5}	10^{-5}
Epochs	100	100
Optimizer	Adam	Adam
LoRA Rank	1	N/A

Table 13: Mean accuracies \pm standard error for 10 test tasks using different retraining and fine-tuning method combinations on ConvAI2.

Algorithm	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	Avg
LoRA-ML	57 \pm 2	41 \pm 4	51 \pm 4	50 \pm 4	51 \pm 4	30 \pm 2	66 \pm 5	47 \pm 4	43 \pm 4	38 \pm 2	47.4 \pm 1.9
SR+LoRA	59 \pm 5	31 \pm 4	50 \pm 7	40 \pm 4	24 \pm 5	20 \pm 2	41 \pm 10	36 \pm 2	23 \pm 5	26 \pm 6	35.0 \pm 4.1
Reptile+LoRA	45 \pm 7	29 \pm 5	35 \pm 6	36 \pm 2	19 \pm 6	21 \pm 3	28 \pm 10	29 \pm 7	21 \pm 5	21 \pm 7	28.4 \pm 5.0
Last-Layer-ML	55 \pm 2	27 \pm 5	51 \pm 5	44 \pm 2	36 \pm 4	25 \pm 4	55 \pm 4	43 \pm 5	33 \pm 4	34 \pm 4	40.2 \pm 1.5
SR+Last-Layer	41 \pm 4	9 \pm 5	29 \pm 3	36 \pm 5	28 \pm 8	15 \pm 4	31 \pm 8	24 \pm 5	15 \pm 6	17 \pm 6	24.5 \pm 4.2
Reptile+Last-Layer	35 \pm 6	13 \pm 4	18 \pm 3	29 \pm 3	19 \pm 9	18 \pm 4	15 \pm 9	17 \pm 5	15 \pm 6	16 \pm 4	19.4 \pm 4.3

Table 14: Mean accuracies (averaged over all tasks) \pm standard error for different retraining methods at varying LoRA ranks during fine-tuning.

Algorithm	Rank 1	Rank 4	Rank 8	Rank 16
LoRA-ML	44.7 \pm 0.8	48.6 \pm 1.6	47.4 \pm 1.9	48.2 \pm 1.5
SR+LoRA	36.3 \pm 4.1	35.5 \pm 4.3	35.0 \pm 4.1	37.1 \pm 4.1
Reptile+LoRA	26.6 \pm 5.8	27.7 \pm 5.3	28.4 \pm 5.0	27.8 \pm 5.9

E.2 Language Experiments

We use the ConvAI2 [23] dataset for the language tasks. ConvAI2 comprises conversations between two personas, where each persona is associated with a list of facts that informs their responses. We model learning the dialogue continuations of each persona as a different task. For each continuation we are given 20 candidate continuations, with one option being the correct continuation.

We retrain starting from the RoBERTa-base model [22] along with the RoBERTa tokenizer. We add a sequence classification head and use cross-entropy loss to predict whether a given continuation is in fact the correct response. We label the incorrect continuations as class 0 and the correct continuation as class 1 for training, and to perform inference we select the continuation with the highest predicted probability of being from class 1. We retrain on the ten tasks (personas) with the most data, and fine-tune on the next ten tasks (personas) with the most data. We perform 5 runs for each experiment (both retraining and fine-tuning) and report the means and the standard errors of the accuracies of the fine-tuned model on the validation sets of the fine-tuning tasks. During retraining, we save the model from the epoch which demonstrates the best average performance on the validation samples for the training tasks. We then fine-tune this retrained model and report the test task accuracies from the last fine-tuning epoch.

The hyperparameters we used are listed below in Tables 15 and 16.

Table 15: Retraining Hyperparameters for language experiments

Hyperparameter	Standard Retraining	LoRA-ML	Last-Layer-ML	Reptile
Learning Rate	$5 \cdot 10^{-5}$	$5 \cdot 10^{-5}$	$5 \cdot 10^{-5}$	$5 \cdot 10^{-5}$
Learning Rate Schedule	Linear	Linear	Linear	Linear
Batch Size	8	8	8	8
Epochs	10	10	10	10
Inner Epochs	N/A	N/A	N/A	20
Optimizer	AdamW	AdamW	AdamW	AdamW
LoRA Rank	N/A	8	N/A	N/A
LoRA Dropout	N/A	0.1	N/A	N/A
LoRA Alpha	N/A	16	N/A	N/A
Outer Learning Rate	N/A	N/A	N/A	10^{-4}

E.3 Asset Information

We use the CIFAR-10 [21] and ConvAI2 [23] datasets in our experiments. CIFAR-10 is publicly available but does not specify an explicit license. ConvAI2 is distributed under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. We also use the MLP-Mixer [35] and RoBERTa-Base [22] model architectures, including the pretrained weights for RoBERTa. The official MLP-

Table 16: Fine-Tuning Hyperparameters for language experiments

Hyperparameter	LoRA Fine-Tuning	Last-Layer Fine-tuning
Learning Rate	10^{-4}	10^{-4}
Learning Rate Schedule	Linear	Linear
Batch Size	8	8
Epochs	10	10
Optimizer	AdamW	AdamW
LoRA Rank	8	N/A
LoRA Dropout	.1	N/A
LoRA Alpha	16	N/A

Mixer implementation is licensed under the Apache License 2.0, while RoBERTa-Base, as accessed via Hugging Face, is licensed under the MIT License.

F Theory Notes

F.1 Non-Uniqueness of Global Min for $T = 2$

Consider $T = 2$, $k = 1$, $d = 2$, $\mathbf{A}^* = \mathbf{0}$, and $\mathbf{u}_t^* = \mathbf{e}_t$ for $t = 1, 2$, where \mathbf{e}_t is the t^{th} standard basis vector. Clearly the ground truth perturbations \mathbf{u}_i^* are orthonormal and thus linearly independent. The set of global minima of \mathcal{L}^* are (\mathbf{A}, \mathbf{U}) such that $\mathbf{A} = \frac{1}{T} \sum_{t=1}^T (\mathbf{u}_t^* \mathbf{u}_t^{*\top} - \mathbf{u}_t \mathbf{u}_t^\top)$ and $\mathbf{u}_t \mathbf{u}_t^\top - \mathbf{u}_t^* \mathbf{u}_t^{*\top} - \frac{1}{T} \sum_{s=1}^T (\mathbf{u}_s \mathbf{u}_s^\top - \mathbf{u}_s^* \mathbf{u}_s^{*\top}) = \mathbf{0}$. It is not hard to see that a global minimum follows from any set values of $\mathbf{u}_1, \mathbf{u}_2$ such that $\mathbf{u}_1 \mathbf{u}_1^\top - \mathbf{u}_2 \mathbf{u}_2^\top = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$. When properly parameterized, this system of equations defines a hyperbola where each point corresponds to a global minimum of \mathcal{L}^* .

F.2 Spurious Local Minima

We observe that for $T \geq 3$, for certain tasks $\mathbf{U}^* = (\mathbf{U}_1^*, \mathbf{U}_2^*, \mathbf{U}_3^*)$, it is possible to find points \mathbf{U} that are local minima, but not global minima. To find these points, we sample true tasks \mathbf{U}^* from a normal distribution and use a numerical solver to find zeros of the gradient of the reduced loss

$$\hat{\mathcal{L}}(\mathbf{U}) = \sum_{t=1}^T \left\| \mathbf{U}_t \mathbf{U}_t^\top - \mathbf{U}_t^* \mathbf{U}_t^{*\top} - \frac{1}{T} \sum_{s=1}^T (\mathbf{U}_s \mathbf{U}_s^\top - \mathbf{U}_s^* \mathbf{U}_s^{*\top}) \right\|_F^2.$$

Through the Schur complement argument used to prove Theorem 4, we can see that $\hat{\mathcal{L}}$ has a spurious local minimum only if \mathcal{L} has a spurious local minimum.

Typically, these zeros are close to the global minimum. Occasionally, it is possible to find a point $\hat{\mathbf{U}}$ with gradients close to 0 and with positive definite Hessians. We then confirm that these are close to the spurious local minimum through the following argument.

Consider the function

$$r(\mathbf{U}) = \text{vec}(\mathbf{U} - \hat{\mathbf{U}})^\top \text{vec}(\nabla \hat{\mathcal{L}}(\mathbf{U})).$$

Clearly, there is a minimum of $\hat{\mathcal{L}}$ in the δ -ball of $\hat{\mathbf{U}}$ if $r(\mathbf{U}) > 0$ for all \mathbf{U} on the boundary of the δ -ball. As r is continuous, if for some small enough $\epsilon, \gamma > 0$ if $r(\mathbf{U}) > \gamma > 0$ for all \mathbf{U} on the ϵ -net of the boundary of the δ -ball, then there exists a spurious local minimum in the δ -ball around $\hat{\mathbf{U}}$. Numerically, such points and ϵ, δ , and γ can be found which would imply that spurious local minima exist, barring any errors due to numerical computation. To confirm, we run gradient descent from this point and observe that the loss stays constant in Figure 12




Figure 12: Loss does not decrease near these spurious local minima

G Example Pseudocode for Minimizing (5)

Algorithm 1 Meta-Adapter Training

- 1: **Input:** Tasks $\mathcal{T}_t, t \in [T]$, learning rate η , number of epochs N_e , batches per epoch N_b
 - 2: **Initialize:** Model parameters $\mathbf{W}_0, \boldsymbol{\theta}_0^{(t)}$ for all $t = 1, \dots, T$
 - 3: **for** epoch $e = 1$ to N_e **do**
 - 4: **for** $b = 1, \dots, N_b$ **do**
 - 5: **for** $t = 1, \dots, T$ **do**
 - 6: Load next batch $\beta_{t,b}$ from \mathcal{T}_i
 - 7: Compute gradient $\mathbf{g}^{(t)} = \nabla_{\mathbf{W}, \boldsymbol{\theta}^{(t)}} \left(\sum_{(\mathbf{x}, \mathbf{y}) \in \beta_{t,b}} \mathcal{L}(\Phi_{\text{FT}}(\mathbf{x}; \mathbf{W}, \boldsymbol{\theta}^{(t)}), \mathbf{y}) \right)$
 - 8: Update adapter parameters: $\boldsymbol{\theta}_{e+1}^{(t)} \leftarrow \boldsymbol{\theta}_e^{(t)} - \eta_e \mathbf{g}^{(t)}$
 - 9: **end for**
 - 10: Update base parameters: $\mathbf{W}_{e+1} \leftarrow \mathbf{W}_e - \eta_e \sum_{t=1}^T \mathbf{g}_{\mathbf{W}}^{(t)}$
 - 11: **end for**
 - 12: **end for**
-

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: Our main claims focus on theoretical guarantees for meta-learning infused with LoRA on linear models. We explicitly state our contributions, as well as their scope relating to LoRA adaptations on linear models, in both the abstract and the introduction. Further, our claims that PEFT-based meta-learning during retraining improves performance over standard retraining is supported through our experiments.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: For all results, we state the scope in which they apply. Our theory applies to linear models using LoRA fine-tuning, and each theoretical statement only applies within our stated assumptions. In the conclusion, we mention future work which includes extending our analysis to more complex model classes and different kinds of PEFT methods.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: For all of our theoretical statements, we state the assumptions along with the result statements. The proofs are in Appendix B.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We disclose all training details in the main body and in Appendices D and E. We further release all code in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility.

In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We release all code, and we use the publicly available datasets CIFAR-10 and ConvAI2.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We explain the experimental methodologies in the paper body along with key hyperparameter choices, like the LoRA rank. The full experimental details are in Appendices D and E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Each set of results is accompanied with a measure of uncertainty. For the synthetic data experiments, we plot average results along with a shaded region showing the full range of values for each setting. For the real data experiments, we report mean results along with standard errors across trials.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: These are reported in Appendices D and E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: This work conforms to each criterion.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This work focuses on the theoretical understanding of model preparation for adaptation using parameter-efficient fine-tuning methods such as LoRA. It does not introduce new models, datasets, or applications with direct societal implications.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not release any new data or models.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Each dataset and model architecture we use is properly cited in the main paper body, and the license information is included in the appendices.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We release our code which includes documentation.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not use experiments of this kind.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not use experiments with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We did not use LLMs for any original or core component of this work.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.