

Understanding Model Merging: A Unified Generalization Framework for Heterogeneous Experts

Qinglun Li^{1*}, Anke Tang^{2*}, Miao Zhang^{1†}, Mengzhu Wang⁴, Quanjun Yin¹, Li Shen^{3†}

¹National University of Defense Technology, Changsha, China

²Wuhan University, Wuhan, China

³School of Cyber Science and Technology, Shenzhen Campus of Sun Yat-sen University, China

⁴Hebei University of Technology, Tianjin, China

liqinglun@nudt.edu.cn, anketang@whu.edu.cn, mathshenli@gmail.com

Abstract

Model merging efficiently aggregates capabilities from multiple fine-tuned models into a single one, operating purely in parameter space without original data or expensive re-computation. Despite empirical successes, a unified theory for its effectiveness under heterogeneous finetuning hyperparameters (e.g., varying learning rates, batch sizes) remains missing. Moreover, the lack of hyperparameter transparency in open-source fine-tuned models makes it difficult to predict merged-model performance, leaving practitioners without guidance on how to fine-tune merge-friendly experts. To address those two challenges, we employ L_2 -Stability theory under heterogeneous hyperparameter environments to analyze the generalization of the merged model x_{avg} . This pioneering analysis yields two key contributions: (i) A unified theoretical framework is provided to explain existing merging algorithms, revealing how they optimize specific terms in our bound, thus offering a strong theoretical foundation for empirical observations. (ii) Actionable recommendations are proposed for practitioners to strategically fine-tune expert models, enabling the construction of merge-friendly models within the pretraining-to-finetuning pipeline. Extensive experiments on the ResNet/Vit family across 20/8 visual classification tasks, involving thousands of finetuning models, robustly confirm the impact of different hyperparameters on the generalization of x_{avg} predicted by our theoretical results.

1. Introduction

The paradigm of pre-training followed by fine-tuning has become the dominant paradigm in computer vision [11, 84, 96], leading to an explosion of powerful foundation mod-

els [5, 17, 98] and a vast ecosystem of specialized “expert” models fine-tuned for specific downstream tasks [25]. Although these experts perform well on their respective tasks, selecting a specific expert model for each downstream task is cumbersome. A critical challenge has emerged: *how can we efficiently aggregate the distinct capabilities of multiple expert models into a single, multi-talented model?* Model merging offers a compelling answer. By operating directly in the parameter space, methods like Model Soups [8, 62, 83] and Task Arithmetic [31] can fuse models without needing access to their original training data or incurring the prohibitive costs of re-training from scratch. Moreover, model merge techniques are also crucial for the advancement of current LLMs [22, 68, 81].

Since model merging methods can operate directly in the parameter space without accessing the original training data and are also computationally efficient, researchers have proposed a wide range of effective techniques from different perspectives [31, 90, 92, 93, 95]. However, these methods are often explained heuristically, and there is still no unified theoretical framework to explain why and when they work. This theoretical gap also leads to significant practical challenges. For example, in open-source communities such as *Hugging Face*, practitioners now have access to a large number of pretrained and fine-tuned models [82]. However, the key hyperparameters used during fine-tuning are often unavailable or non-transparent. This lack of transparency introduces two main difficulties: (i) The performance of model merging becomes highly unpredictable [93], and (ii) The **try-all-and-error** process to identify suitable combinations greatly increases experimental cost [29, 97]. Furthermore, without theoretical guidance, practitioners attempting to merge multiple “expert models” face two central questions: (i) Since the underlying mechanisms that determine the effectiveness of different model merging methods remain unclear, it is difficult to select an appropriate merging

*These authors contributed equally to this work.

†Corresponding authors

strategy for a given set of models; and (ii) Without knowing how fine-tuning processes and their hyperparameters influence the final merging performance, it is challenging to fine-tune new models that can serve as good candidates for future merging—so-called *merge-friendly* experts (i.e., expert models with good merging compatibility) [19, 85]. This leaves the community with an unclear theoretical framework to guide the full pretrain-to-merge pipeline.

To bridge these critical gaps in both theory and practice, our work seeks to answer two fundamental questions:

- (1) *What makes different model merging algorithms effective, and can these methods be unified within a single framework?*
- (2) *How do the hyperparameters of these algorithms affect their performance, and how can practitioners adjust them to improve the performance of x_{avg} ?*

To answer these questions, we provide a rigorous theoretical analysis of model merging through the lens of L_2 -Stability [44, 46]. Our primary contribution is the first excess error bound for models merged from experts trained in heterogeneous hyperparameter environments. This pioneering theoretical result yields two key insights. First, it provides a *unified framework* that explains the effectiveness of diverse merging algorithms. We theoretically demonstrate how different methods, from simple averaging [31] to sparsity-based approaches [95], can be interpreted as strategies to optimize distinct terms within our bound, thus providing a strong theoretical foundation for previously empirical observations. Second, our analysis yields *actionable recommendations* for practitioners. By elucidating the precise role each hyperparameter (e.g., learning rate η , batch size b_i , training steps K_i) plays in the generalization of the final merged model, we offer a strategic guide for constructing merge-friendly experts.

Empirically, we validate our theory through extensive experiments, performing thousands of fine-tuning and merging trials on the ResNet family across diverse image classification tasks [71, 76]. The results demonstrate a remarkable alignment between our theory and empirical generalization trends. Moreover, this work establishes the first unified and predictive mathematical framework for model merging, transforming it from an empirical craft into a principled and systematic discipline, and offering theory-grounded insights to guide future developments in the field.

Table 1 summarizes the main challenges and our contributions.

2. Related Works

This section reviews diverse *model merging methods* and the current *theoretical advancement* in the field. At the end

of each subsection, we explicitly outline the remaining gaps and open questions, thereby clearly positioning the contributions of this work.

2.1. Model Merging Methods

Model merging techniques can be broadly categorized by the stage at which they intervene as pre- and during-merging methods, following the taxonomy of Yang et al. [91].

Pre-Merging Methods focus on preparing models to be more compatible for merging, primarily by addressing the geometric misalignment of their parameter spaces. A key strategy is *Weight Alignment* [1, 23, 34, 38, 47, 51, 54, 65, 89], which seeks to find permutation symmetries or other geometric relationships to align models within the same loss basin before averaging. Git Re-basin [1] and its extensions [60] exemplify this approach. Another tactic is to modify the fine-tuning procedure itself, for instance, through *Linearized Fine-tuning* in the tangent space of a pretrained model, which helps disentangle task-specific knowledge and reduce interference upon merging [48, 57].

During-Merging Methods represent the core fusion algorithms. These methods span from basic methods to advanced strategies that mitigate task interference, representing one of the most active areas of current research. These methods include: (i) *Basic methods*: The simplest approaches include direct weight averaging, as popularized by Model Soups [3, 10, 39, 62, 83, 88, 99, 100], and Task Arithmetic [31, 35, 67], which adds or subtracts vectors representing specific task capabilities. While foundational, their performance degrades when models exhibit conflicting parameter updates. (ii) *Weighted-based merging*: To improve upon simple averaging, these methods assign adaptive weights to different models or their components. Some methods learn these coefficients using an auxiliary dataset [93], while others rely on heuristics like the Fisher information matrix to determine parameter importance [16, 52, 58]. (iii) *Subspace-based merging*: A highly effective strategy for reducing interference is first to sparsify models and merge them in a shared subspace [27, 69, 94]. TIES-Merging [90] mitigates parameter conflicts by zeroing inconsistent signs and keeping high-magnitude updates, while DARE [95] later showed that randomly pruning and rescaling task vectors can further enhance performance. and (iv) *Post-calibration methods*: A more recent direction aims to calibrate the merged model after fusion. This is motivated by the observation that even a successful merge in weight space can lead to a “representation bias”, where the merged model’s internal representations differ from those of the expert models [73, 74, 92]. Representation Surgery [92] addresses this by introducing a lightweight module to correct the merged model’s feature space in a post-hoc step.

Open Questions: Model merging has developed a rich set of empirically successful techniques, yet the underlying

Table 1. Summary of challenges in model merging and our contributions.

Aspect	Prevailing Challenges & Gaps	Our Contribution
Theoretical	<ul style="list-style-type: none"> - Lack of a unified theory to explain various methods. - Effectiveness is explained by heuristics, not principles. - Unclear why merging works in heterogeneous settings. 	<p>A Unified Generalization Framework: We derive the first excess error bound via L_2-Stability, explaining <i>why</i> different methods are effective.</p>
Practical	<ul style="list-style-type: none"> - Merged model performance is highly unpredictable. - No clear guidance on how to fine-tune models for merging. - High trial-and-error costs for practitioners. 	<p>Actionable Guidance for Practitioners: Our bound predicts hyperparameter impact, providing a guide to create ‘merge-friendly’ experts.</p>
Empirical	<ul style="list-style-type: none"> - Lack of systematic, large-scale studies connecting theory to practice under varied hyperparameters. 	<p>Extensive Empirical Validation: We finetune thousands of models across 20 tasks, showing our theory strongly aligns with real-world results.</p>

principles unifying their success remain unclear. Currently, no foundational theory explains, within a single framework, why diverse strategies—such as weight alignment before merging, sparsification during merging, or representation correction after merging—lead to **deteriorated** merged mod-els. This work addresses that gap by offering a unified theoretical perspective grounded in generalization theory.

2.2. Theoretical Advancement in Model Merging

Theoretical understanding of model merging significantly lags its empirical progress. As cataloged in Yang et al. [91], existing theoretical work can be grouped by the context in which models are merged.

First, a robust line of research analyzes the averaging of model checkpoints along a single training trajectory, including methods such as SWA [32, 78, 79], EMA [6, 53], LAWA [2, 36], and SeWA [80]. The success of these methods is generally attributed to their ability to converge to wide, flat minima in the loss landscape, which are known to promote better generalization [26, 77]. However, such analyses are confined to single-task settings and thus do not extend to our multi-task scenario.

Second, some studies explain the merging of models fine-tuned on the *same dataset* but with different initializations or hyperparameters. The dominant theory here is *Linear Mode Connectivity (LMC)* [18, 20, 21]. LMC posits that the solutions found by SGD for overparameterized networks are not isolated points but are connected by linear paths of low loss. This provides a powerful geometric intuition for why models trained on the same data distribution can be effectively averaged, especially after alignment [1].

Third, and most relevant to our paper, is the theory for merging models fine-tuned on *different datasets or tasks*. This is the most challenging and least understood scenario. The few existing analyses are often tailored to specific methods. For example, Ortiz-Jimenez et al. [57] use the Neural Tangent Kernel (NTK) to provide a compelling theoretical link between task arithmetic and the spectral properties of the model, but this analysis is specific to their proposed linearized fine-tuning method and is not applicable to

heterogeneous hyperparameter settings of this paper.

Open Questions: The current theoretical landscape remains fragmented, lacking a unified framework that explains the generalization of merged models trained on diverse tasks under heterogeneous hyperparameter settings. Theories like LMC do not fully account for the data heterogeneity between tasks, and method-specific analyses do not provide a universal perspective. Therefore, a critical gap remains: We lack a theoretical framework that can both unify different merging methods and directly relate fine-tuning hyperparameters to the final generalization error of the merged model. This paper aims to fill this critical gap by leveraging L_2 -Stability to derive the first such unifying generalization bound.

3. Problem Formulation

We consider N tasks, each finetuned from a pretrained model with parameters \mathbf{x}_0 . The goal of model merging is to independently train N expert models on different tasks and find a merged parameter $\mathbf{x}_{avg} = \mathbf{x}_0 + \sum_{i=1}^N \lambda_i (\mathbf{x}_i - \mathbf{x}_0)$ that performs well across all N tasks [31, 93], where \mathbf{x}_i denotes the expert model parameters obtained by training \mathbf{x}_0 on task i . Specifically, model merging aims to minimize the following global population risk:

$$F(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{z \sim \mathcal{P}_i} [\ell(\mathbf{x}; z)], \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ and \mathcal{P}_i denotes the data distribution for task i , and the distributions differ across tasks. The loss function $\ell(\mathbf{x}; z) : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}^+$ is non-negative, and z is sampled from the sample space \mathcal{Z} . Since the distributions \mathcal{P}_i are typically unknown, the global population risk $F(\mathbf{x})$ cannot be computed exactly. Instead, we can access a training dataset $\mathcal{D} = \cup_{i=1}^N \mathcal{D}_i$, where $\mathcal{D}_i = \{z_j^{(i)}\}_{j=1}^{n_i}$ is the training set for task i of size n_i . We then approximate $F(\mathbf{x})$ using the global empirical risk:

$$f(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{x}), \quad (2)$$

where $f_i(\mathbf{x}) = \frac{1}{n_i} \sum_{j=1}^{n_i} \ell(\mathbf{x}, z_j^{(i)})$ is the empirical risk for task i , and the samples $z_j^{(i)} \stackrel{\text{i.i.d.}}{\sim} \mathcal{P}_i$ are drawn i.i.d. from \mathcal{P}_i .

To address the two questions raised in Section 1, we need to theoretically establish the relationship between the generalization bound of \mathbf{x}_{avg} and the relevant hyperparameters. Next, we introduce several definitions for analyzing generalization bounds and explain their practical significance.

Excess Error[44, 45]. In this work, we focus on analyzing the excess error of algorithm, defined as $\mathcal{E}(\mathbf{x}) := F(\mathbf{x}) - f(\hat{\mathbf{x}})$, where $\hat{\mathbf{x}}$ denotes the empirical risk minimizer (ERM). Moreover, the excess error can be decomposed as

$$\mathcal{E}(\mathbf{x}) = \underbrace{F(\mathbf{x}) - f(\mathbf{x})}_{\mathcal{E}_G : \text{generalization gap}} + \underbrace{f(\mathbf{x}) - f(\hat{\mathbf{x}})}_{\mathcal{E}_O : \text{optimization error}}. \quad (3)$$

Where the \mathcal{E}_G represents the generalization error caused by approximating the unknown data distributions $\mathcal{P}_i, \forall i \in [N]$. The second term \mathcal{E}_O captures the optimization error on the training data \mathcal{D} .

Remark 1 (Understanding Excess Error): Excess error measures the gap between the risk on the true data distribution $F(\mathbf{x})$ and the empirical risk on the training dataset $f(\hat{\mathbf{x}})$. It can always be decomposed into the sum of the optimization error \mathcal{E}_O and the generalization error \mathcal{E}_G . Specifically, the optimization error quantifies the gap between the solution $f(\mathbf{x})$ found by our algorithm on the training set and the theoretically optimal solution $f(\hat{\mathbf{x}})$, reflecting the capability and efficiency of the optimization algorithm itself. The generalization error measures the gap between the solution $f(\mathbf{x})$ found on the training set and the optimal solution $F(\mathbf{x})$ on the true data distribution, meaning how well the patterns learned from the training data generalize to unseen data. By combining both components, excess error provides a more comprehensive measure of an algorithm's optimization efficiency and the quality of its solution.

Definition 1 (Perturbed datasets). Let the global dataset be denoted by $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$ and $\mathcal{D}_i = \{z_1, \dots, z_{n_i}\}$, where each \mathcal{D}_i represents the local dataset of the i -th task with $|\mathcal{D}_i| = n_i$ for all $i \in [N]$. Consider another global dataset $\tilde{\mathcal{D}} = \bigcup_{i=1}^N \tilde{\mathcal{D}}_i$ and $\tilde{\mathcal{D}}_i = \{\tilde{z}_1, \dots, \tilde{z}_{n_i}\}$, independently sampled from \mathcal{Z} such that $z_j, \tilde{z}_j \sim \mathcal{P}_i$ whenever $z_j \in \mathcal{D}_i$. For j , which is randomly selected from the indices $1, \dots, n_i$, define $\mathcal{D}^{(i)} \triangleq \{z_1, \dots, z_{j-1}, \tilde{z}_j, z_{j+1}, \dots, z_{n_i}\}$ as the perturbed version of \mathcal{D} , obtained by randomly replacing the j -th element with \tilde{z}_j , and define $\tilde{\mathcal{D}} \triangleq \bigcup_{j \neq i} \mathcal{D}_j \cup \mathcal{D}^{(i)}$.

Definition 2 (l_2 on-average model stability) Let \mathcal{A} be a randomized algorithm. By leveraging the notation for perturbed datasets introduced in Definition 1, we say that \mathcal{A} is l_2 on-average model ε -stable if the following inequality

holds:

$$\mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N \|\mathcal{A}(\mathcal{D}) - \mathcal{A}(\mathcal{D}^{(i)})\|_2^2 \right] \leq \varepsilon^2. \quad (4)$$

The expectation \mathbb{E} is taken over the random draws of the datasets \mathcal{D} and $\tilde{\mathcal{D}}$, as well as the internal randomness of the algorithm \mathcal{A} .

Remark 2 (Advantages of L_2 -Stability): Compared with Uniform Stability [49, 66], L_2 -Stability [44, 46] offers the following advantages: (1). Weaker assumptions and better practical applicability. Specifically, L_2 -Stability removes the bounded-gradient assumption required by Uniform Stability. This represents a significant theoretical improvement, as in deep learning, there exist loss functions with unbounded gradients—for example, the mean squared error. Hence, the assumption behind L_2 -Stability broadens its theoretical coverage and enhances its practical relevance. (2). Greater robustness to individual sample perturbations. Uniform Stability measures the worst-case upper bound, so if a single data point experiences a large perturbation, the bound becomes overly large. In contrast, L_2 -Stability measures an average-case bound, making it more robust to variations in individual samples.

4. Theoretical Analysis

In this section, we first state the assumptions required for the theoretical analysis, then analyze the excess error $\mathcal{E}(\mathbf{x}_{avg})$ of model merge methods. Based on the derived excess-error bounds, we explain why different model merge methods are effective, how various hyperparameters influence their performance, and provide practical guidelines for improving empirical results.

Assumption 1 (L -smoothness). The loss function ℓ is L -smooth i.e. $\exists L > 0$ such that $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d, z \in \mathcal{D}$, $\|\nabla \ell(\mathbf{x}; z) - \nabla \ell(\mathbf{y}; z)\|_2 \leq L \|\mathbf{x} - \mathbf{y}\|_2$.

Assumption 2 (Bounded Variance). The local stochastic gradient $\mathbf{g}_i = \frac{1}{|\xi_i|} \sum_{z \in \xi_i} \nabla \ell(\mathbf{x}; z)$ is unbiased $\mathbb{E}[\mathbf{g}_i] = \nabla f_i(\mathbf{x})$ for all $\xi_i \in \mathcal{D}_i, i \in [N]$. Let $b_i = |\xi_i|$ denote the batch size used for task i and there exists $\sigma_i^2 > 0$ such that $\mathbb{E}\|\mathbf{g}_i - \nabla f_i(\mathbf{x})\|^2 \leq \frac{\sigma_i^2}{b_i}$, for all $\xi_i \in \mathcal{D}_i, i \in [N]$.

Assumption 3 (Bounded Heterogeneity). There exists $\zeta_i^2 > 0$ such that $\mathbb{E}\|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq \zeta_i^2$, for any $i \in [N]$ and $\mathbf{x} \in \mathbb{R}^d$.

Assumption 4 (Coefficients) Let $\lambda_1, \dots, \lambda_N$ be the coefficients used to scale task vectors in the model merge algorithm. We assume that $\lambda_i \geq 0$ and $\sum_{i=1}^N \lambda_i = 1$.

Remark 3 (Rationale for Different Assumptions): For Assumption 2 and 3, unlike standard assumptions for algorithmic convergence, we adopt gradient variance and heterogeneity assumptions that are dependent on the index i , which aligns with the practical training hyperparameter settings of each expert model.

Notably, unlike analyses based on Uniform-Stability [45, 49, 66], our assumptions do not include the bounded-gradient assumption. Next, we present the generalization error bound and excess error bound based on L_2 -Stability.

4.1. Excess Error of Model Merging Method

Before presenting the upper bound of Excess Error, we first introduce an important lemma that establishes the connection between L_2 -Stability and Generalization Error. The proofs of all subsequent theorems and lemmas can be found in the Appendix A.

Lemma 1 (Generalization via on-average model stability)

Let $\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{D}^{(i)}$ be constructed as Definition 1. Let $\gamma > 0$. If for any z , the function $f(\mathbf{x}; z)$ is nonnegative and L -smooth, then

$$\begin{aligned} \mathcal{E}_G &\leq \frac{1}{2\gamma} \mathbb{E}_{\mathcal{D}, \mathcal{A}} [\|\nabla f(\mathcal{A}(\mathcal{D}))\|^2] \\ &+ \frac{L + \gamma}{2} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} [\|\mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D})\|^2]. \end{aligned}$$

where $\mathcal{A}(\cdot)$ denotes a learning algorithm, and we denote $\mathbf{x} = \mathcal{A}(\mathcal{D})$ as the model produced by applying \mathcal{A} to the dataset \mathcal{D} . From Lemma 1, we know that the generalization error \mathcal{E}_G on the left-hand side is controlled by the upper bound of the gradient norm and L_2 -Stability on the right-hand side. Therefore, we only need to derive the upper bound of the gradient norm and the L_2 -Stability bound separately to control \mathcal{E}_G . Moreover, by combining a common assumption $\mathcal{E}_O(\mathbf{x}_{avg}) \leq C \cdot \mathbb{E} \|\nabla f(\mathbf{x}_{avg})\|^2$ (This assumption will be detailed in Appendix A, Step 4.), we obtain the following stability-based expression for the excess error:

$$\begin{aligned} \mathcal{E}(\mathbf{x}_{avg}) &\leq \frac{L + \gamma}{2} \underbrace{\mathbb{E} \|\mathbf{x}_{avg} - \tilde{\mathbf{x}}_{avg}\|^2}_{I_1: \text{Model Stability}} \\ &+ \left(\frac{1}{2\gamma} + C \right) \underbrace{\mathbb{E} \|\nabla f(\mathbf{x}_{avg})\|^2}_{I_2: \text{Gradient Norm}}. \end{aligned} \quad (5)$$

From inequality (5), we know that by controlling the upper bounds of the first term (*Model Stability*) and the second term (*Gradient Norm*) separately, we can obtain an upper bound on the excess error. In the following, we present the upper bounds for each term.

(I). Upper Bound of I_1 : The proof of I_1 proceeds in two steps: first, derive $\tilde{I}_1 \triangleq \|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2$; second, obtain an upper bound for I_1 . Since $\mathbf{x}_{avg} = \sum_{i=1}^N \lambda_i \mathbf{x}_i^{K_i}$, $\tilde{\mathbf{x}}_{avg} = \sum_{i=1}^N \lambda_i \tilde{\mathbf{x}}_i^{K_i}$ always holds, bounding I_1 requires first bounding \tilde{I}_1 . The following theorem provides an upper bound for I_1 .

Theorem 1 (Upper Bound of Model Stability): Let $\mathcal{D}, \tilde{\mathcal{D}}$ be constructed as in Definition 1. Let \mathbf{x}_{avg} and $\tilde{\mathbf{x}}_{avg}$ denote the model parameters obtained by finetuning on \mathcal{D} and $\tilde{\mathcal{D}}$, respectively, and then applying the model-merge method. Then we have the following:

$$\mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}} \|\mathbf{x}_{avg} - \tilde{\mathbf{x}}_{avg}\|^2 \leq 16\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2 \right)$$

(II). Upper Bound of I_2 : In real-world open-source model parameters, models \mathbf{x}_i with different capabilities are often obtained using different training hyperparameters, such as different optimizers or varying numbers of training steps. Therefore, to make our results more consistent with practical scenarios, when bounding I_2 , we introduce both *training heterogeneity* (i.e., task-specific hyperparameters such as batch size b_i , number of steps K_i , gradient noise σ_i , etc., which differ across tasks) and *data heterogeneity* (i.e., distributional divergence parameters ζ_i between tasks i and j). This incorporation ensures that our results possess greater generality and broader applicability.

Theorem 2 (Upper Bound of Gradient Norm) Under Assumptions 1, 2, let $\bar{K} = \frac{1}{N} \sum_{i=1}^N K_i$, $\eta_l = \sqrt{\frac{N}{\bar{K}}}$ and all task use the SGD optimizer. Then, the gradient of the surrogate function $\tilde{f}(\mathbf{x}) = \sum_{i=1}^N \lambda_i f_i(\mathbf{x})$ at the averaged model \mathbf{x}_{avg} is bounded as follows:

$$\begin{aligned} \mathbb{E} [\|\nabla \tilde{f}(\mathbf{x}_{avg})\|^2] &\leq \frac{4(\tilde{f}(\mathbf{x}_0) - \tilde{f}_{\text{inf}})}{\sqrt{N\bar{K}}} + \frac{4L\sigma^2 A_1}{\sqrt{N\bar{K}}} \\ &+ \frac{6NL^2\sigma^2 A_2}{\bar{K}} + \frac{5NL^2\zeta^2 A_3}{\bar{K}}. \end{aligned} \quad (6)$$

Where $A_1 = N \sum_{i=1}^N \frac{\bar{K}}{K_i} \lambda_i^2$, $A_2 = \sum_{i=1}^N \lambda_i (K_i - 1)$, $\tilde{f}_{\text{inf}} = \min_{\mathbf{x}} \tilde{f}(\mathbf{x})$, $\sigma = \max_i \{\sigma_i\}$, $\zeta = \max_i \{\zeta_i\}$ and $A_3 = \max_i \{K_i (K_i - 1)\}$.

Furthermore, let the right-hand side of inequality (6) be denoted as ϵ_{sgd} . Then, we have the following inequality, which ensures that even under heterogeneous conditions, the original objective function $f(\mathbf{x})$ can still converge, although its bound is not as tight as that in inequality 14:

$$\mathbb{E} [\|\nabla f(\mathbf{x}_{avg})\|^2] \leq 2[\chi_{\mathbf{p}|\lambda}^2 + 1] \epsilon_{sgd} + 2\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 \quad (7)$$

where $\chi_{\mathbf{p}||\lambda}^2 = \sum_{i=1}^N \frac{(\frac{1}{N} - \lambda_i)^2}{\lambda_i^2}$ is the chi-square divergence between the weight coefficient vectors $\lambda = [\lambda_1, \lambda_2, \dots, \lambda_N] \in \mathbb{R}^N$ and $\mathbf{p} = [\frac{1}{N}, \frac{1}{N}, \dots, \frac{1}{N}] \in \mathbb{R}^N$.

(III). Upper Bound of Excess Error: By combining the bounds of I_1 and I_2 with inequality (5), we can readily obtain the upper bound of the excess error.

Theorem 3 (Tight Upper Bound of Excess Error) By combining the conditions and results from Theorems 1 and 2, and choosing an appropriate γ^* in inequality (5), we obtain the following tighter excess error bound.

$$\begin{aligned} \mathcal{E}(\mathbf{x}_{avg}) \leq & 8(L+1)\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2 \right) \\ & + (2C+1) \left[\chi_{\mathbf{p}||\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}||\lambda}^2 + 1) \epsilon_{sgd} \right] \end{aligned}$$

Remark 4 (Interpretation of Classical Model Merge Methods via the Excess Error Bound): The excess error bound presented in Theorem 3 provides a unified theoretical framework for understanding the mechanisms behind various model merging methods. Each method can be interpreted as modifying different terms in Theorem 3 to balance the optimization error \mathcal{E}_O and the generalization error \mathcal{E}_G , thereby minimizing the excess error. Next, we will use the bound in Theorem 3 to analyze the reasons behind the effectiveness of each mainstream model-merge method. This provides an answer to Question (1) in Section 1.

- **Pre-Merging Methods (e.g., Git Re-basin [1]):** These methods focus on aligning the parameter spaces of expert models before they are averaged. From our theoretical perspective, this alignment serves to reduce the initial task heterogeneity ζ_i^2 . By finding permutation symmetries that place models in a common basin, these techniques ensure that the subsequent merging step operates on a set of models that are already more geometrically and functionally concordant, thereby tightening our excess error bound from the outset. Moreover, our Stability Bound scales with the Lipschitz constant L (representing loss landscape curvature). [43] (SAM) minimizes the dominant eigenvalue of the Hessian ($\lambda_{\max}(\mathbf{H})$), which effectively reduces the local Lipschitz constant L , thereby directly shrinking the stability penalty. Similarly, [70] (Linearization) approximates the complex loss landscape with a tangent plane. This structural constraint enforces near-zero curvature (Hessian ≈ 0), minimizing the non-linear deviations that cause instability, thus strictly tightening the stability bound.
- **During-Merging Methods:** These algorithms directly manipulate the fusion process itself, and our framework reveals the distinct trade-offs each strategy makes. (i) **Basic Methods (Simple Averaging [83], Task Arithmetic [31]):**

This approach, which sets $\lambda_i = 1/N$, greedily optimizes one term in our bound: it sets the chi-square divergence $\chi_{\mathbf{p}||\lambda}^2$ to zero, completely eliminating the penalty from optimizing a surrogate objective. However, its agnosticism to task heterogeneity makes it vulnerable; outlier tasks with large ζ_i^2 can inflate both the stability and optimization error components, leading to suboptimal performance. This motivates the pruning strategy proposed in the TIES method [90]. (ii) **Weighted-based Merging (e.g., AdaMerging [93]):** These methods perform a data-driven optimization of our entire error bound. They learn non-uniform coefficients λ that strategically accept a non-zero surrogate penalty ($\chi_{\mathbf{p}||\lambda}^2 > 0$) in exchange for a greater reduction in the heterogeneity-dependent terms (e.g., $\sum \lambda_i \zeta_i^2$). This explicitly down-weights outlier tasks to achieve a better overall balance between the competing terms in our bound. (iii) **Subspace-based Merging (e.g., TIES [90], DARE [95]):** These techniques directly reduce task heterogeneity ζ_i^2 at a fine-grained, parameter level. Since $f = \sum \lambda_i f_i$, ∇f and $\mathbb{E} \|\nabla f_i - \nabla f\|^2 \leq \zeta_i^2$. Thus, ζ_i^2 mathematically represents the variance of gradients across tasks. By eliminating conflicts in update directions, pruning methods (e.g., TIES, DARE) perform statistical trimming of outlier gradients, reducing variance and tightening the ζ_i^2 , which, according to our bound in Theorem 3, directly tightens both the stability and optimization error components to improve generalization. Other methods, such as [30] and [50], follow the same mechanism as TIES and likewise reduce ζ_i^2 . and (iv) **Post-calibration methods (e.g., Representation Surgery [92]):** These methods act as a corrective step after the main fusion is complete. From our framework's perspective, they implicitly acknowledge that the merged model \mathbf{x}_{avg} may still possess a non-trivial representation error. Instead of further modifying the parameters, they introduce a lightweight, task-specific module to compensate for this residual error in the representation space, effectively calibrating the model's outputs without altering the core merged parameters.

Remark 5 (The Impact of Different Hyperparameters on Excess Error and Practical Fine-tuning Recommendations): Our theoretical framework not only explains the differences between merging algorithms but also elucidates the critical role of various hyperparameters in the fine-tune-then-merge process. In the Section 5.2, we provide a detailed analysis of how each hyperparameter influences the final generalization bound. Moreover, practical fine-tuning recommendations for practitioners are presented in Appendix A.1.2. This addresses Question (1) in Section 1.

Remark 6 (Generalization to Adam.) While Theorem 6 focuses on SGD for clarity, Theorem 12 (Appendix) establishes a general framework. As discussed in [75] and our Appendix (Eq.(11)), adaptive optimizers (Adam) can be

modeled by modifying the effective step-size vector \mathbf{a}_i (i.e., adaptive learning rates). This alters the coefficients in the bound but preserves the fundamental stability-optimization trade-off, ensuring the theory remains applicable.

5. Experiments

In this section, we empirically validate the theoretical insights from Section 4. Our goal is to show that the proposed generalization bound accurately predicts the behavior of the merged model x_{avg} under different finetuning hyperparameters. We analyze the effects of finetuning steps (K_i), batch size (b_i), learning rate (η_l), data ratio (α_i), and number of merged tasks (N). Our code is publicly available at https://gitcode.com/tanganke/stability_model_merging.

5.1. Experimental Setup

Datasets. Our experimental testbed consists of 20 diverse visual classification tasks, ranging from standard benchmarks like CIFAR-10 and SVHN to more complex, fine-grained challenges such as Stanford Cars and Food-101, as well as scene recognition with SUN397 and satellite imagery with EuroSAT. The complete dataset can be found in the Appendix B. This variety ensures a thorough and varied evaluation of model merging performance. For each experiment, we evaluate each merged model on the joint test set—the union of all task test sets—providing a unified measure of its multi-task performance.

Models and Merging Methods. We use three widely adopted backbones from the ResNet family—ResNet-18, ResNet-50, and ResNet-152—as well as the CLIP model based on the Vision Transformer (ViT) architecture. Each expert model is initialized from the same pretrained checkpoint, obtained from Hugging Face, and then finetuned on its respective task. During fine-tuning, we use SGD as the optimizer for ResNet-based models and Adam for ViT-based models. Unless otherwise specified as the variable under investigation, we use a standard set of hyperparameters for finetuning. For ResNet-based backbones, we adopt Simple Averaging, which is equivalent to Task Arithmetic [31] when the scaling factor is $1/N$. For ViT-based frameworks, we evaluate more advanced merging methods, including TIES [90], DARE [95], and AdaMerging [93], to validate our theory.

Evaluation Protocol. To ensure robust and statistically significant results, we employ a randomized sampling procedure. For each experimental setting, we randomly sample 15 distinct groups of tasks from our pool of 20. We then perform the merging process for each group and report the *average loss and accuracy* across these 15 trials. All experiments are conducted separately for each of the three ResNet backbones. For experiments based on the Transformer architecture, we evaluate on 8 tasks [93]. Unlike the ResNet

setting, we fix the number of tasks to 8 in all experiments except those that explicitly study the effect of task count.

5.2. Verifying Theoretical Predictions

We now present the results of five targeted experiments, each designed to validate a specific aspect of our theoretical framework. The results are visualized in Figure 1&2.

5.2.1. Impact of Fine-tuning Steps (K_i)

Theoretical Prediction. As analyzed in Appendix A.1.2, the number of finetuning steps, K_i , presents a fundamental trade-off. Our bound in Theorem 3 indicates that increasing K_i reduces the optimization error component (specifically, the ϵ_{sgd} term, which decays with \bar{K}), but inflates the model stability component, which grows linearly with K_i . This suggests an optimal number of steps that balances underfitting (large optimization error) and over-specialization (high instability). Consequently, the performance of the merged model depends on which of these competing factors dominates. We might observe performance first increasing and then decreasing with K_i , forming an inverted U-shape. Alternatively, if the stability cost grows faster than the optimization benefit from the start, performance could also decrease monotonically.

Setup and Results. We merge models finetuned for $K_i \in \{1000, 2000, 3000, 4000\}$ steps, while holding other hyperparameters constant. As depicted in Figure 2(a), the empirical results show that the loss exhibits a monotonic increase or decrease, with accuracy changing inversely, which aligns with our theoretical predictions. Different models have distinct optimal values K_i^* . When the actual number of steps exceeds K_i^* , the loss rises because the increase in model instability dominates the overall excess error. Conversely, when the steps are fewer than K_i^* , the loss decreases as the reduction in optimization error becomes the dominant factor. These observations strongly support our theoretical prediction of the optimization–generalization trade-off governed by K_i .

5.2.2. Impact of Batch Size (b_i)

Theoretical Prediction. Our theory highlights a dual role for the batch size b_i . A larger b_i reduces stochastic gradient variance, improving optimization, but Theorem 1 shows that the stability term grows linearly with b_i (i.e., $3b_i\zeta_i^2/n_i$), reflecting increased sensitivity to data perturbations under task heterogeneity. Therefore, b_i represents a trade-off between optimization and generalization. However, from Theorem 3, we observe that the coefficient of b_i in the generalization error bound is η_l^2 . Since $\eta_l = 0.001$ in our experiments, the overall effect of b_i is primarily reflected through its influence on the optimization error ϵ_{sgd} . Consequently, as b_i increases, the merged model tends to exhibit better generalization performance.

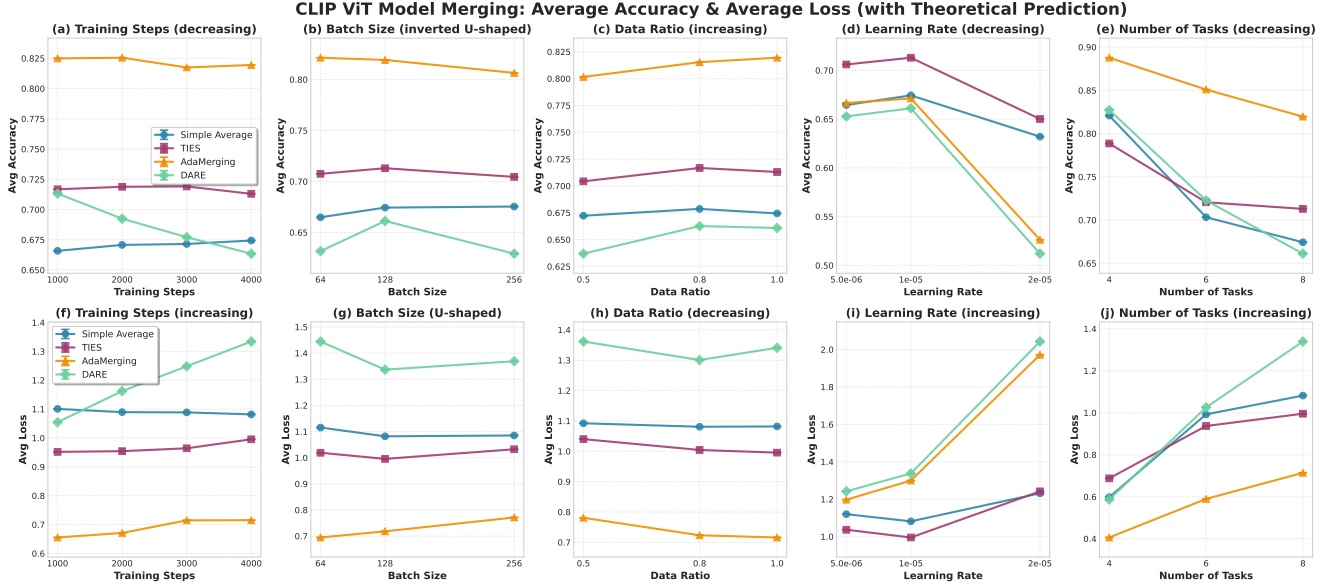


Figure 1. **Empirical Validation of Theoretical Predictions in Transformer.** The experiments evaluate four representative methods under different algorithmic hyperparameters in terms of average accuracy and average loss, both measured on x_{avg} . Each subfigure also includes the theoretically predicted trend from Theorem 3. As can be observed, all methods closely follow the theoretical predictions, demonstrating the applicability of our framework.

Comprehensive Analysis of Model Merging Stability Across Different Hyperparameters (with Theoretical Prediction)

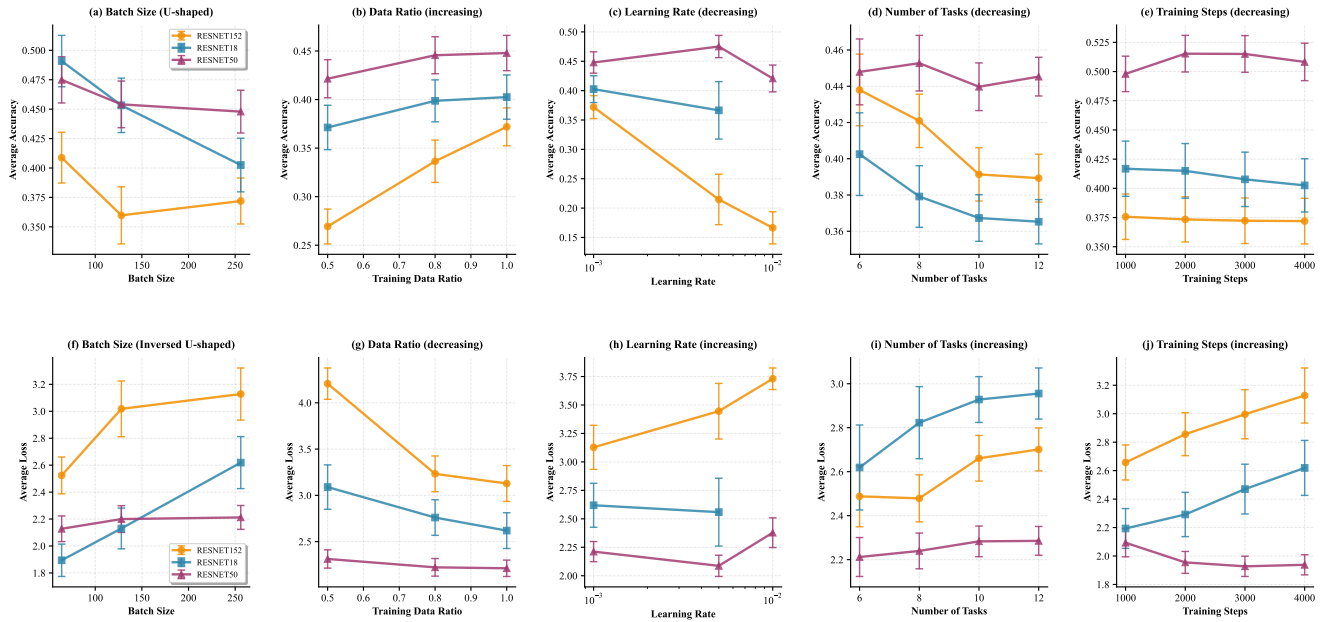


Figure 2. **Empirical Validation of Theoretical Predictions in ResNet.** We analyze the impact of key finetuning hyperparameters on the performance of the merged model (x_{avg}). Each subplot corresponds to varying one hyperparameter while keeping others fixed. Performance is measured by accuracy (higher is better) and loss (lower is better) on the joint test set, averaged over 15 random task groups. The observed trends robustly align with the trade-offs predicted by our L_2 -Stability-based generalization error bound in Theorem 3.

Setup and Results. We vary the batch size $b_i \in \{64, 128, 256\}$ for all finetuning processes. The results, depicted in Figure 1&2, show a consistent trend of improved

performance as the batch size increases, with accuracy rising and loss decreasing monotonically. This empirical finding provides compelling validation for our theoretical anal-

ysis. As predicted in our theoretical discussion, under our experimental setting with a small learning rate ($\eta_l = 0.001$ in ResNet and $\eta_l = 0.00001$ in ViT), the negative impact on model stability is heavily suppressed by the η_l^2 coefficient. Consequently, the dominant effect is the reduction in optimization error (ϵ_{sgd}) afforded by the more accurate gradient estimates from larger batches. This result underscores the precision of our bound, which correctly predicts that while a trade-off between optimization and generalization exists in principle, its empirical manifestation is conditioned by other critical hyperparameters like the learning rate.

5.2.3. Impact of Learning Rate (η_l)

Theoretical Prediction. The learning rate, η_l , is identified by our theory as a dominant factor controlling stability from Theorem 3. The entire stability component in Theorem 3 is scaled by η_l^2 (where η_l is a function of the finetuning learning rate η_l). Our framework thus predicts that smaller learning rates are paramount for ensuring the stability required for effective model merging, while larger rates can cause the stability term to explode, leading to poor generalization regardless of the optimization progress.

Setup and Results. We finetune expert models using learning rates $\eta_l \in \{0.001, 0.005, 0.01\}$ in ResNet and $\{1e-5, 5e-6, 2e-5\}$ in ViT. The empirical results in Figure 1&2 provide compelling evidence for this prediction. We observe a sharp and consistent degradation in performance (both higher loss and lower accuracy) as the learning rate increases. In fact, when $\eta_l = 0.01$, we observed that all merged models fine-tuned on ResNet-18 collapsed. This validates that maintaining a small learning rate during finetuning is critical not only for preserving pretrained features but, more fundamentally, for ensuring the algorithmic stability necessary for the merged model to generalize well.

5.2.4. Impact of Data Ratio (α_i)

Theoretical Prediction. The effect of the per-task dataset size, n_i (which our data ratio α_i directly controls), is uniquely beneficial and unambiguous according to our theory. We refer to our final excess error bound in Theorem 3, which is composed of a *Model Stability* component and an *Optimization Error* component. The dataset size n_i appears exclusively in the denominator of the entire stability term, which is scaled by $\sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2 \right)$. Consequently, increasing n_i strictly and monotonically tightens this crucial part of the bound.

Setup and Results. We simulate varying dataset sizes by using a fraction $\alpha_i \in \{0.5, 0.8, 1.0\}$ of the training data for each task. As shown in Figure 1&2, the performance of the merged model improves monotonically as the data ratio increases. This result aligns with our theory and statistical learning principles, confirming that fine-tuning on more data produces more stable, generalizable, and merge-friendly expert models.

5.2.5. Impact of Number of Tasks (N)

Theoretical Prediction. The number of tasks, N , introduces competing effects in our excess error bound, but an analysis of their scaling rates and coefficients predicts a clear negative trend. On one hand, the optimization error term, ϵ_{sgd} , provides a benefit that improves at a rate of $O(1/\sqrt{N})$ as N increases. However, we argue this positive effect is minor in practice. Its coefficient, $4(\tilde{f}(\mathbf{x}_0) - \tilde{f}_{\text{inf}})$, is small because finetuning begins from a strong pretrained model \mathbf{x}_0 that is already close to the optimum \tilde{f}_{inf} . Furthermore, the negative impact from task heterogeneity accumulates much more quickly. Terms in our bound related to heterogeneity grow at least linearly with N . Crucially, their coefficients, which are tied to inter-task variance, are not small. This creates a fundamental imbalance: a minor, slowly-improving benefit is easily outweighed by a significant, fast-growing penalty. Therefore, we predict that as N increases, the penalty from accumulating task heterogeneity will dominate, leading to a consistent decline in the merged model’s generalization performance.

Setup and Results. We vary the number of merged tasks $N \in \{6, 8, 10, 12\}$. In Figure 1&2(e), we observe that performance generally degrades as more tasks are merged using simple averaging. This suggests that for our set of diverse tasks, the negative impact of increasing heterogeneity outweighs the potential optimization benefits. This result highlights the importance of the heterogeneity terms in our bound and implicitly motivates more advanced merging techniques (as discussed in Remark 4) that can adaptively handle a large number of diverse tasks.

6. Conclusion

This paper first establishes an L_2 -stability-based framework that provides a unified theoretical explanation for various model merging methods. Building upon this framework: (i) We clarify why different merging algorithms are effective, showing that each method tightens the excess error bound from a different perspective; and (ii) We analyze how fine-tuning hyperparameters affect the generalization of \mathbf{x}_{avg} and offer practical guidance for training merge-friendly expert models. These contributions bridge the gap between theory and practice in model merging. In future work, we will relax key assumptions (e.g., non-smooth settings) to generalize our theory. Meanwhile, we are expected to inspire the development of more effective merging methods guided by our theoretical bounds. For example, the proposed theoretical bounds can guide the development of an AutoML approach.

References

- [1] Samuel K Ainsworth, Jonathan Hayase, and Siddhartha S Srinivasa. Git re-basin: Merging models modulo permuta-

- tion symmetries. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 3, 6, 9
- [2] Niccolò Ajroldi, Antonio Orvieto, and Jonas Geiping. When, Where and Why to Average Weights? In *Forty-Second International Conference on Machine Learning*, 2025. 3
- [3] Benjamin Biggs, Arjun Seshadri, Yang Zou, Achin Jain, Aditya Golatkar, Yusheng Xie, Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. Diffusion Soup: Model Merging for Text-to-Image Diffusion Models, 2024. 2
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, pages 446–461. Springer, 2014. 11
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners, 2020. 1
- [6] Dan Busbridge, Jason Ramapuram, Pierre Ablin, Tatiana Likhomanenko, Eeshan Gunesh Dhekane, Xavier Suau Quadros, and Russell Webb. How to Scale Your EMA. *Advances in Neural Information Processing Systems*, 36: 73122–73174, 2023. 3
- [7] Zachary Charles and Dimitris Papailiopoulos. Stability and generalization of learning algorithms that converge to global optima. In *International conference on machine learning*, pages 745–754. PMLR, 2018. 8
- [8] Atoosa Chegini, Hamid Kazemi, Seyed Iman Mirzadeh, Dong Yin, Maxwell Horton, Moin Nabi, Mehrdad Farajtabar, and Keivan Alizadeh. Model Soup for Better RLHF: Weight Space Averaging to Improve Alignment in LLMs. In *NeurIPS 2024 Workshop on Fine-Tuning in Modern Machine Learning: Principles and Scalability*, 2024. 1
- [9] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 11
- [10] Alexandra Chronopoulou, Matthew Peters, Alexander Fraser, and Jesse Dodge. AdapterSoup: Weight Averaging to Improve Generalization of Pretrained Language Models. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2054–2063, Dubrovnik, Croatia, 2023. Association for Computational Linguistics. 2
- [11] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling Instruction-Finetuned Language Models. *Journal of Machine Learning Research*, 25(70):1–53, 2024. 1
- [12] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. 11
- [13] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep learning for classical japanese literature. *arXiv preprint arXiv:1812.01718*, 2018. 11
- [14] Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011. 11
- [15] Gregory Cohen, Saeed Afshar, Jonathan Tapon, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE, 2017. 11
- [16] Donato Crisostomi, Alessandro Zirilli, Antonio Andrea Gargiulo, Maria Sofia Bucarelli, Simone Scardapane, Fabrizio Silvestri, Iacopo Masi, and Emanuele Rodolà. MASS: MoErging through Adaptive Subspace Selection, 2025. 2
- [17] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [18] Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*, pages 1309–1318. PMLR, 2018. 3
- [19] Sebastian Dziaozio, Vishaal Udandarao, Karsten Roth, Ameya Prabhu, Zeynep Akata, Samuel Albanie, and Matthias Bethge. How to Merge Your Multimodal Models Over Time? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20479–20491, 2025. 2
- [20] Rahim Entezari, Hanie Sedghi, Olga Saukh, and Behnam Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. In *The Tenth International Conference on Learning Representations*, 2022. 3
- [21] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pages 3259–3269. PMLR, 2020. 3
- [22] Zichuan Fu, Xian Wu, Yejing Wang, Wanyu Wang, Shanshan Ye, Hongzhi Yin, Yi Chang, Yefeng Zheng, and Xiangyu Zhao. Training-free llm merging for multi-task learning. *arXiv preprint arXiv:2506.12379*, 2025. 1
- [23] Antonio Andrea Gargiulo, Donato Crisostomi, Maria Sofia Bucarelli, Simone Scardapane, Fabrizio Silvestri, and Emanuele Rodolà. Task Singular Vectors: Reducing Task Interference in Model Merging. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18695–18705, 2025. 2
- [24] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *Neural information processing: 20th international conference, ICONIP 2013, daegu, korea, november 3-7, 2013. Proceedings, Part III 20*, pages 117–124. Springer, 2013. 11
- [25] Yunhui Guo, Honghui Shi, Abhishek Kumar, Kristen Grauman, Tajana Rosing, and Rogerio Feris. Spottune: transfer learning through adaptive fine-tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4805–4814, 2019. 1
- [26] Moritz Hardt, Benjamin Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In *International conference on machine learning*, pages 1225–1234. PMLR, 2016. 3
- [27] Yifei He, Yuzheng Hu, Yong Lin, Tong Zhang, and Han Zhao. Localize-and-Stitch: Efficient Model Merging via Sparse Task Arithmetic. *Transactions on Machine Learning Research*, 2024. 2
- [28] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019. 11
- [29] Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. LoraHub: Efficient Cross-Task Generalization via Dynamic LoRA Composition. In *First Conference on Language Modeling*, 2024. 1
- [30] Chenyu Huang, Peng Ye, Tao Chen, Tong He, Xiangyu Yue, and Wanli Ouyang. Emr-merging: Tuning-free high-performance model merging, 2024. 6
- [31] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *ICLR*, 2023. 1, 2, 3, 6, 7, 4, 8
- [32] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *Conference on Uncertainty in Artificial Intelligence*, pages 876–885. PMLR, 2018. 3
- [33] Prateek Jain, Purushottam Kar, et al. Non-convex optimization for machine learning. *Foundations and Trends® in Machine Learning*, 10(3-4):142–363, 2017. 7
- [34] Dong-Hwan Jang, Sangdoo Yun, and Dongyoon Han. Model Stock: All we need is just a few fine-tuned models, 2024. 2
- [35] Ruochen Jin, Bojian Hou, Jiancong Xiao, Weijie J. Su, and Li Shen. Fine-Tuning Attention Modules Only: Enhancing Weight Disentanglement in Task Arithmetic. In *The Thirteenth International Conference on Learning Representations*, 2024. 2
- [36] Jean Kaddour. Stop Wasting My Time! Saving Days of ImageNet and BERT Training with Latest Weight Averaging. In *Has It Trained Yet? NeurIPS 2022 Workshop*, 2022. 3
- [37] Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 795–811. Springer, 2016. 8
- [38] Edan Kinderman, Itay Hubara, Haggai Maron, and Daniel Soudry. Foldable SuperNets: Scalable Merging of Transformers with Different Initializations and Tasks, 2024. 2
- [39] Anat Kleiman, Gintare Karolina Dziugaite, Jonathan Frankle, Sham Kakade, and Mansheej Paul. Soup to go: Mitigating forgetting during continual learning with model averaging, 2025. 2
- [40] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. 11
- [41] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 11
- [42] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 11
- [43] Yeoreum Lee, Jinwook Jung, and Sungyong Baik. Mitigating parameter interference in model merging via sharpness-aware fine-tuning. *arXiv preprint arXiv:2504.14662*, 2025. 6
- [44] Yunwen Lei and Yiming Ying. Fine-grained analysis of stability and generalization for stochastic gradient descent. In *International Conference on Machine Learning*, pages 5809–5819. PMLR, 2020. 2, 4
- [45] Qinglun Li, Miao Zhang, Yingqi Liu, Qunjun Yin, Li Shen, and Xiaochun Cao. Boosting the performance of decentralized federated learning via catalyst acceleration. *arXiv preprint arXiv:2410.07272*, 2024. 4, 5
- [46] Qinglun Li, Yingqi Liu, Miao Zhang, Xiaochun Cao, Qunjun Yin, and Li Shen. Unveiling the power of multiple gossip steps: A stability-based generalization analysis in decentralized training. *arXiv preprint arXiv:2510.07980*, 2025. 2, 4
- [47] Zexi Li, Zhiqi Li, Jie Lin, Tao Shen, Tao Lin, and Chao Wu. Training-time neuron alignment through permutation subspace for improving linear mode connectivity and model fusion. *arXiv preprint arXiv:2402.01342*, 2024. 2
- [48] Tian Yu Liu and Stefano Soatto. Tangent Model Composition for Ensembling and Continual Fine-tuning. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18630–18640, Paris, France, 2023. IEEE. 2
- [49] Yingqi Liu, Qinglun Li, Jie Tan, Yifan Shi, Li Shen, and Xiaochun Cao. Understanding the stability-based generalization of personalized federated learning. In *The Thirteenth International Conference on Learning Representations*, 2025. 4, 5
- [50] Zhenyi Lu, Chenghao Fan, Wei Wei, Xiaoye Qu, Dangyang Chen, and Yu Cheng. Twin-merging: Dynamic integration of modular expertise in model merging, 2024. 6
- [51] Daniel Marczak, Simone Magistri, Sebastian Cygert, Bartłomiej Twardowski, Andrew D. Bagdanov, and Joost

- van de Weijer. No Task Left Behind: Isotropic Model Merging with Common and Task-Specific Subspaces. In *Forty-Second International Conference on Machine Learning*, 2025. 2
- [52] Michael Matena and Colin Raffel. Merging models with fisher-weighted averaging. In *Advances in Neural Information Processing Systems*, pages 17703–17716, 2022. 2
- [53] Daniel Morales-Brotons, Thijs Vogels, and Hadrien Hendrikx. Exponential Moving Average of Weights in Deep Learning: Dynamics and Benefits, 2024. 3
- [54] Anshul Nasery, Jonathan Hayase, Pang Wei Koh, and Sewoong Oh. PLeaS - Merging Models with Permutations and Least Squares. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 30493–30502, 2025. 2
- [55] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, page 4. Granada, 2011. 11
- [56] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008. 11
- [57] Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *Advances in Neural Information Processing Systems*, 36:66727–66754, 2023. 2, 3
- [58] Oleksiy Ostapenko, Zhan Su, Edoardo Maria Ponti, Laurent Charlin, Nicolas Le Roux, Matheus Pereira, Lucas Caccia, and Alessandro Sordoni. Towards Modular LLMs by Building and Reusing a Library of LoRAs, 2024. 2
- [59] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. 11
- [60] Fidel A Guerrero Peña, Heitor Rapela Medeiros, Thomas Dubail, Masih Aminbeidokhti, Eric Granger, and Marco Pedersoli. Re-basin via implicit sinkhorn differentiation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20237–20246, 2023. 2
- [61] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 11
- [62] Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor, Laure Soulier, and Matthieu Cord. Rewarded soups: Towards Pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards. In *Thirty-Seventh Conference on Neural Information Processing Systems*, 2023. 1, 2
- [63] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013. 11
- [64] Johannes Stalkamp, Marc Schlipf, Jan Salmen, and Christian Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32:323–332, 2012. 11
- [65] George Stoica, Daniel Bolya, Jakob Brandt Björner, Pratik Ramesh, Taylor Hearn, and Judy Hoffman. ZipIt! Merging Models from Different Tasks without Training. In *The Twelfth International Conference on Learning Representations*, 2023. 2
- [66] Tao Sun, Dongsheng Li, and Bao Wang. Stability and generalization of decentralized stochastic gradient descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9756–9764, 2021. 4, 5
- [67] Wenju Sun, Qingyong Li, Wen Wang, Yangliao Geng, and Boyang Li. Task Arithmetic in Trust Region: A Training-Free Model Merging Approach to Navigate Knowledge Conflicts. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 5178–5187, New York, NY, USA, 2025. Association for Computing Machinery. 2
- [68] Derek Tam, Margaret Li, Prateek Yadav, Rickard Brühl Gabrielsson, Jiacheng Zhu, Kristjan Greenewald, Mikhail Yurochkin, Mohit Bansal, Colin Raffel, and Leshem Choshen. Llm merging: Building llms efficiently through merging. In *NeurIPS 2024 Competition Track*, 2024. 1
- [69] Anke Tang, Li Shen, Yong Luo, Liang Ding, Han Hu, Bo Du, and Dacheng Tao. Concrete Subspace Learning based Interference Elimination for Multi-task Model Fusion, 2023. 2
- [70] Anke Tang, Li Shen, Yong Luo, Yibing Zhan, Han Hu, Bo Du, Yixin Chen, and Dacheng Tao. Parameter efficient multi-task model fusion with partial linearization. *arXiv preprint arXiv:2310.04742*, 2023. 6
- [71] Anke Tang, Enneng Yang, Li Shen, Yong Luo, Han Hu, Bo Du, and Dacheng Tao. Merging models on the fly without retraining: A sequential approach to scalable continual model merging. *arXiv preprint arXiv:2501.09522*, 2025. 2
- [72] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11*, pages 210–218. Springer, 2018. 11
- [73] Fanqi Wan, Xinting Huang, Deng Cai, Xiaojun Quan, Wei Bi, and Shuming Shi. Knowledge Fusion of Large Language Models. In *The Twelfth International Conference on Learning Representations*, 2023. 2
- [74] Fanqi Wan, Longguang Zhong, Ziyi Yang, Ruijun Chen, and Xiaojun Quan. FuseChat: Knowledge Fusion of Chat Models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 21629–21653, Suzhou, China, 2025. Association for Computational Linguistics. 2
- [75] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances*

- in neural information processing systems*, 33:7611–7623, 2020. 6, 5, 7
- [76] Ke Wang, Nikolaos Dimitriadis, Guillermo Ortiz-Jimenez, François Fleuret, and Pascal Frossard. Localizing Task Information for Improved Model Merging and Compression. In *Proceedings of the 41st International Conference on Machine Learning*, pages 50268–50287. PMLR, 2024. 2
- [77] Peng Wang, Li Shen, Zerui Tao, Shuaida He, and Dacheng Tao. Generalization analysis of stochastic weight averaging with general sampling. In *International Conference on Machine Learning*, 2024. 3
- [78] Peng Wang, Li Shen, Zerui Tao, Yan Sun, Guodong Zheng, and Dacheng Tao. A unified analysis for finite weight averaging. *arXiv preprint arXiv:2411.13169*, 2024. 3
- [79] Peng Wang, Shengchao Hu, Zerui Tao, Guoxia Wang, Dianhai Yu, Li Shen, Quan Zheng, and Dacheng Tao. Sewa: Selective weight average via probabilistic masking. *arXiv preprint arXiv:2502.10119*, 2025. 3
- [80] Peng Wang, Shengchao Hu, Zerui Tao, Guoxia Wang, Dianhai Yu, Li Shen, Quan Zheng, and Dacheng Tao. SeWA: Selective Weight Average via Probabilistic Masking, 2025. 3
- [81] Zhixiang Wang, Zhenyu Mao, Yixuan Qiao, Yunfang Wu, and Biye Li. Optimal brain iterative merging: Mitigating interference in llm merging. *arXiv preprint arXiv:2502.12217*, 2025. 1
- [82] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. HuggingFace’s Transformers: State-of-the-art Natural Language Processing, 2020. 1
- [83] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR, 2022. 1, 2, 6
- [84] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. Robust fine-tuning of zero-shot models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7949–7961, New Orleans, LA, USA, 2022. IEEE. 1
- [85] Chengyue Wu, Teng Wang, Yixiao Ge, Zeyu Lu, Ruisong Zhou, Ying Shan, and Ping Luo. π -Tuning: Transferring Multimodal Foundation Models with Optimal Multi-task Interpolation. In *Proceedings of the 40th International Conference on Machine Learning*, pages 37713–37727. PMLR, 2023. 2
- [86] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashionmnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 11
- [87] Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010. 11
- [88] Guofu Xie, Xiao Zhang, Ting Yao, and Yunsheng Shi. Bone Soups: A Seek-and-Soup Model Merging Approach for Controllable Multi-Objective Generation, 2025. 2
- [89] Yichu Xu, Xin-Chun Li, Le Gan, and De-Chuan Zhan. Weight scope alignment: A frustratingly easy method for model merging. *arXiv preprint arXiv:2408.12237*, 2024. 2
- [90] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36:7093–7115, 2023. 1, 2, 6, 7, 9
- [91] Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang, Xiaochun Cao, Jie Zhang, and Dacheng Tao. Model merging in llms, mllms, and beyond: Methods, theories, applications and opportunities. *arXiv preprint arXiv:2408.07666*, 2024. 2, 3
- [92] Enneng Yang, Li Shen, Zhenyi Wang, Guibing Guo, Xiaojun Chen, Xingwei Wang, and Dacheng Tao. Representation surgery for multi-task model merging. *arXiv preprint arXiv:2402.02705*, 2024. 1, 2, 6, 9
- [93] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. *ICLR*, 2024. 1, 2, 3, 6, 7, 4, 9
- [94] Xin Yi, Shunfan Zheng, Linlin Wang, Xiaoling Wang, and Liang He. A safety realignment framework via subspace-oriented model fusion for large language models. *Knowledge-Based Systems*, 306:112701, 2024. 2
- [95] Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*, 2024. 1, 2, 6, 7, 9
- [96] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021. 1
- [97] Chujie Zheng, Ziqi Wang, Heng Ji, Minlie Huang, and Nanyun Peng. Model Extrapolation Expedites Alignment. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1025–1041, Vienna, Austria, 2025. Association for Computational Linguistics. 1
- [98] Hongling Zheng, Li Shen, Anke Tang, Yong Luo, Han Hu, Bo Du, Yonggang Wen, and Dacheng Tao. Learning from models beyond fine-tuning. *Nature Machine Intelligence*, 7(1):6–17, 2025. 1
- [99] Yuyan Zhou, Liang Song, Bingning Wang, and Weipeng Chen. MetaGPT: Merging Large Language Models Using Model Exclusive Task Arithmetic. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Lan-*

guage Processing, pages 1711–1724, Miami, Florida, USA, 2024. Association for Computational Linguistics. [2](#)

- [100] Max Zimmer, Christoph Spiegel, and Sebastian Pokutta. Sparse Model Soups: A Recipe for Improved Pruning via Model Averaging, 2023. [2](#)

Understanding Model Merging: A Unified Generalization Framework for Heterogeneous Experts

Supplementary Material

A Appendix	1
A.1. Preliminary Lemmas	1
B Experiment Details	11
B.1. Datasets Setup	11
C More Results	12
C.1. The Average Performance of Individual Models Before Merging	12
C.2. Single-Task Finetuning Performance	12
C.3. Detailed Experimental Results	14

A. Appendix

In this part, we provide the supplementary materials to prove the main theorem.

A.1. Preliminary Lemmas

Then, given a learning algorithm $\mathcal{A}(\cdot)$ and a training dataset \mathcal{D} , we denote $\mathbf{x} = \mathcal{A}(\mathcal{D})$ as the model generated by the method \mathcal{A} with given data. Then, the upper bound of the model's generalization gap is established in the following theorem:

Lemma 2 (Generalization via on-average model stability) *Let $\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{D}^{(i)}$ be constructed as Definition 1. Let $\gamma > 0$. If for any z , the function $f(\mathbf{x}; z)$ is nonnegative and L -smooth, then*

$$\mathbb{E}_{\mathcal{D}, \mathcal{A}}[F(\mathcal{A}(\mathcal{D})) - f(\mathcal{A}(\mathcal{D}))] \leq \frac{1}{2\gamma} \mathbb{E}_{\mathcal{D}, \mathcal{A}}[\|\nabla f(\mathcal{A}(\mathcal{D}))\|^2] + \frac{L + \gamma}{2} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}}[\|\mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D})\|^2].$$

Proof 1 *According to the definition of generalization error, we have:*

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}, \mathcal{A}}[F(\mathcal{A}(\mathcal{D})) - f(\mathcal{A}(\mathcal{D}))] \\ &= \mathbb{E}_{\mathcal{D}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N (F(\mathcal{A}(\mathcal{D}))) \right] - \mathbb{E}_{\mathcal{D}, \mathcal{A}} [f(\mathcal{A}(\mathcal{D}))] \\ &= \mathbb{E}_{\mathcal{D}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\tilde{z}_i} [\ell(\mathcal{A}(\mathcal{D}), \tilde{z}_i)] \right] - \mathbb{E}_{\mathcal{D}, \mathcal{A}} [f(\mathcal{A}(\mathcal{D}))] \\ &= \mathbb{E}_{\mathcal{D}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\tilde{z}_i} [\ell(\mathcal{A}(\mathcal{D}), \tilde{z}_i)] \right] - \mathbb{E}_{\mathcal{D}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N \ell(\mathcal{A}(\mathcal{D}); z_i) \right] \\ &= \mathbb{E}_{\mathcal{D}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\tilde{\mathcal{D}}} [\ell(\mathcal{A}(\mathcal{D}), \tilde{z}_i)] \right] - \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} [\ell(\mathcal{A}(\mathcal{D}^{(i)}); \tilde{z}_i)] \\ &= \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \left[\frac{1}{N} \sum_{i=1}^N \left(\ell(\mathcal{A}(\mathcal{D}), \tilde{z}_i) - \ell(\mathcal{A}(\mathcal{D}^{(i)}); \tilde{z}_i) \right) \right]. \end{aligned}$$

Since $\ell(\mathbf{x}; z)$ is L -smooth, i.e., $\forall \mathbf{x}, \mathbf{y} \in \mathcal{X}, \forall z \in \mathcal{Z}, \|\nabla \ell(\mathbf{x}; z) - \nabla \ell(\mathbf{y}; z)\| \leq L\|\mathbf{x} - \mathbf{y}\|$, and by the definition of empirical risk $f(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \ell(\mathbf{x}; z_j)$, we have

$$\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| = \left\| \frac{1}{N} \sum_{i=1}^N \nabla \ell(\mathbf{x}; z_j) - \frac{1}{N} \sum_{i=1}^N \nabla \ell(\mathbf{y}; z_j) \right\|$$

$$\begin{aligned}
&\leq \frac{1}{N} \sum_{i=1}^N \|\nabla \ell(\mathbf{x}; z_j) - \nabla \ell(\mathbf{y}; z_j)\| \\
&\leq L \|\mathbf{x} - \mathbf{y}\|
\end{aligned}$$

Therefore, the empirical risk is also L -smooth, and we have the following inequality.

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}, \mathcal{A}}[F(\mathcal{A}(\mathcal{D})) - f(\mathcal{A}(\mathcal{D}))] &\leq \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \left[\langle \mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D}), \frac{1}{N} \sum_{i=1}^N \nabla \ell(\mathcal{A}(\mathcal{D}); \tilde{z}_j) \rangle \right] \\
&\quad + \frac{L}{2} \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \left[\|\mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D})\|^2 \right]
\end{aligned}$$

Using Cauchy's inequality, for any $\gamma > 0$, we have

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \left[\langle \mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D}), \frac{1}{N} \sum_{i=1}^N \nabla \ell(\mathcal{A}(\mathcal{D}); \tilde{z}_i) \rangle \right] &\leq \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \|\mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D})\| \cdot \|\nabla f(\mathcal{A}(\mathcal{D}))\| \\
&\leq \frac{\gamma}{2} \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \|\mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D})\|^2 + \frac{1}{2\gamma} \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} \|\nabla f(\mathcal{A}(\mathcal{D}))\|^2
\end{aligned}$$

By further simplification, we obtain the final result.

$$\mathbb{E}_{\mathcal{D}, \mathcal{A}}[F(\mathcal{A}(\mathcal{D})) - f(\mathcal{A}(\mathcal{D}))] \leq \underbrace{\frac{L+\gamma}{2} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathcal{D}, \tilde{\mathcal{D}}, \mathcal{A}} [\|\mathcal{A}(\mathcal{D}^{(i)}) - \mathcal{A}(\mathcal{D})\|^2]}_{\text{On-average Model stability}} + \frac{1}{2\gamma} \underbrace{\mathbb{E}_{\mathcal{D}, \mathcal{A}} \|\nabla f(\mathcal{A}(\mathcal{D}))\|^2}_{\text{Gradient norm}}$$

In fact, we modified the proof of the lemma from [44], replacing the $f(\mathcal{A}(\mathcal{D}))$ on the right-hand side with a gradient $\mathbb{E}_{\mathcal{A}, \mathcal{D}}[\|\nabla f(\mathcal{A}(\mathcal{D}))\|^2]$. This adjustment better captures the impact of data heterogeneity on the generalization error. With this lemma, obtaining the desired generalization bound reduces to controlling the l_2 on-average model stability of the decentralized algorithm A .

Step 1: Proof of Local Model Stability:

Lemma 3 Letting non-negative objectives $f_i, \forall i \in [N]$ satisfy L -smoothness and Assumptions 2, 3. We suppose the i -th task preserves dataset \mathcal{D}_i and $|\mathcal{D}_i| = n_i$ samples. Its perturbed dataset $\mathcal{D}_i^{(j)}$ has the perturbed sample $\tilde{z}_j \in \mathcal{D}_i$ with probability 1. If the i -th task conducts K_i mini-batch SGD steps with batch-size b_i , and non-increasing learning rate $\eta_t = \Theta\left(\frac{1}{KL}\right)$, we prove the stability of local iteration on the i -th task as

$$\mathbb{E}_{\mathcal{D}_i, \mathcal{D}_i'} \|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 \leq 16K\eta_i^2 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i\zeta_i^2}{n_i} \right) \quad (8)$$

Proof 2 Noting that the dataset \mathcal{D}_i and the perturbed dataset $\tilde{\mathcal{D}}_i$ differ by at most one sample, there are two possible cases to consider when running local mini-batch size SGD.

First Case: In the first case, the local mini-batch size SGD does not select the perturbed sample from either \mathcal{D}_i or $\tilde{\mathcal{D}}_i$, in which case we have

$$\begin{aligned}
&\mathbb{E}[\|\mathbf{x}_i^{k+1} - \tilde{\mathbf{x}}_i^{k+1}\|^2 | \tilde{z} \notin \xi] \\
&\leq \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k - \eta_l(\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k)\|^2 \\
&\leq \mathbb{E} \left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2 \|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 - 2\eta_l \langle \mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k, \mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k \rangle \right] \\
&\leq \mathbb{E} \left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2 \|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 - 2\eta_l \langle \mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k, \nabla f_i(\mathbf{x}_i^k) - \nabla f_i(\tilde{\mathbf{x}}_i^k) \rangle \right] \\
&\leq \mathbb{E} \left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2 \|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 + 2\eta_l L \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 \right]
\end{aligned}$$

$$\begin{aligned}
&\leq (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\mathbb{E}\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 \\
&= (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\mathbb{E}\|\mathbf{g}_i^k \pm \nabla f_i(\mathbf{x}_i^k) - \tilde{\mathbf{g}}_i^k \pm \nabla f_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\stackrel{(i)}{\leq} (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + 2\eta_l^2\mathbb{E}\|\mathbf{g}_i^k - \nabla f_i(\mathbf{x}_i^k) - \tilde{\mathbf{g}}_i^k + \nabla f_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\quad + 2\eta_l^2\mathbb{E}\|\nabla f_i(\mathbf{x}_i^k) - \nabla f_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\leq (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + 2\eta_l^2\mathbb{E}\|\mathbf{g}_i^k - \nabla f_i(\mathbf{x}_i^k)\|^2 + 2\eta_l^2\mathbb{E}\|\tilde{\mathbf{g}}_i^k - \nabla f_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\quad + 2\eta_l^2\mathbb{E}\|\nabla f_i(\mathbf{x}_i^k) - \nabla f_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\leq (1 + 2\eta_l L + 2\eta_l^2 L^2)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2\sigma_i^2}{b_i}
\end{aligned}$$

where we note that the stochastic gradients \mathbf{g}_i^k and $\tilde{\mathbf{g}}_i^k$ only differ in model parameters (\mathbf{x}_i^k and $\tilde{\mathbf{x}}_i^k$) and the mini batch data samples are identical in (i).

Second Case: In the second case, local mini-batch SGD samples a batch of data that involves the perturbed sample \tilde{z} from \mathcal{D}_i and \mathcal{D}'_i , which happens with probability b_i/n_i . Analogously, we have

$$\begin{aligned}
&\mathbb{E}[\|\mathbf{x}_i^{k+1} - \tilde{\mathbf{x}}_i^{k+1}\|^2 | \tilde{z} \in \xi] \\
&\leq \mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k - \eta_l(\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k)\|^2 \\
&\leq \mathbb{E}\left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 - 2\eta_l\langle \mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k, \mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k \rangle\right] \\
&\leq \mathbb{E}\left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 - 2\eta_l\langle \mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k, \nabla f_i(\mathbf{x}_i^k) - \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k) \rangle\right] \\
&\leq \mathbb{E}\left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 \right. \\
&\quad \left. - 2\eta_l\langle \mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k, \frac{1}{n_i} \sum_{j=1}^{n_i} (\nabla \ell(\mathbf{x}_i^k; z_j) - \nabla \ell(\tilde{\mathbf{x}}_i^k; z_j)) \rangle + \nabla \ell(\tilde{\mathbf{x}}_i^k; z) - \nabla \ell(\tilde{\mathbf{x}}_i^k; \tilde{z})\right] \\
&\stackrel{(i)}{\leq} \mathbb{E}\left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 - 2\eta_l\langle \mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k, \nabla f_i(\mathbf{x}_i^k) - \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k) \rangle\right] \\
&\leq \mathbb{E}\left[\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 + 2\eta_l L\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2\right] \\
&\leq (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\mathbb{E}\|\mathbf{g}_i^k - \tilde{\mathbf{g}}_i^k\|^2 \\
&\stackrel{(ii)}{\leq} (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \eta_l^2\mathbb{E}\|\mathbf{g}_i^k \pm \nabla f_i(\mathbf{x}_i^k) - \tilde{\mathbf{g}}_i^k \pm \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\leq (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2\sigma_i^2}{b_i} + 2\eta_l^2\mathbb{E}\|\nabla f_i(\mathbf{x}_i^k) - \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k)\|^2
\end{aligned} \tag{9}$$

where (i) utilizes the fact that $z, \tilde{z} \stackrel{i.i.d.}{\sim} \mathcal{P}_i$, implying that $\mathbb{E}_{z, \tilde{z}} [\nabla \ell(\tilde{\mathbf{x}}_i^k; z) - \nabla \ell(\tilde{\mathbf{x}}_i^k; \tilde{z})] = 0$, due to the expectation being taken over the local data distribution. Since the perturbed samples z and \tilde{z} are independently and identically distributed (i.i.d.), and under Assumption 1, which characterizes the smoothness of the local data distribution, the above holds. Moreover, Assumption 2 is assumed to hold for SGD with batch size 1, i.e., $\|\nabla \ell(\mathbf{x}; z) - \nabla f_i(\mathbf{x})\| \leq \sigma_i^2, \forall \mathbf{x} \in \mathcal{X}, z \in \mathcal{D}_i \sim \mathcal{P}_i$. We define the virtual local objective as $\tilde{f}_i(\tilde{\mathbf{x}}_i^k) := \frac{1}{|\mathcal{D}'_i|} \sum_{z \in \mathcal{D}'_i} \ell(\tilde{\mathbf{x}}_i^k; z)$, which differs from $f_i(\tilde{\mathbf{x}}_i^k) = \frac{1}{|\mathcal{D}_i|} \sum_{z \in \mathcal{D}_i} \ell(\tilde{\mathbf{x}}_i^k; z)$ by considering a perturbed sample. In step (ii), we use the fact that both $\tilde{f}_i(\tilde{\mathbf{x}}_i^k)$ and the corresponding stochastic gradient $\tilde{\mathbf{g}}_i^k$ satisfy Assumption 2. In addition, $\tilde{f}_i(\mathbf{x})$ also satisfies Assumption 3 regarding its alignment with the global objective f , in the same way as $f_i(\mathbf{x})$ does, for all $\mathbf{x} \in \mathcal{X}$. This holds since z and \tilde{z} are drawn i.i.d. from \mathcal{P}_i , for every $i \in [N]$. Based on the above discussion, we conclude that

$$\begin{aligned}
&\mathbb{E}[\|\mathbf{x}_i^{k+1} - \tilde{\mathbf{x}}_i^{k+1}\|^2 | \tilde{z} \in \xi] \\
&\leq (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2\sigma_i^2}{b_i} + 2\eta_l^2\mathbb{E}\|\nabla f_i(\mathbf{x}_i^k) - \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\leq (1 + 2\eta_l L)\mathbb{E}\|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2\sigma_i^2}{b_i} + 2\eta_l^2\mathbb{E}\|\nabla f_i(\mathbf{x}_i^k) \pm \nabla f(\mathbf{x}_i^k) - \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k) \pm \nabla f(\tilde{\mathbf{x}}_i^k)\|^2
\end{aligned}$$

$$\begin{aligned}
&\leq (1 + 2\eta_l L) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2 \sigma_i^2}{b_i} \\
&\quad + 2\eta_l^2 \mathbb{E} \|\nabla f_i(\mathbf{x}_i^k) - \nabla f(\mathbf{x}_i^k) + \nabla f(\tilde{\mathbf{x}}_i^k) - \nabla \tilde{f}_i(\tilde{\mathbf{x}}_i^k) + \nabla f(\mathbf{x}_i^k) - \nabla f(\tilde{\mathbf{x}}_i^k)\|^2 \\
&\leq (1 + 2\eta_l L) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2 \sigma_i^2}{b_i} + 12\eta_l^2 \zeta_i^2 + 6\eta_l^2 L^2 \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 \\
&= (1 + 2\eta_l L + 6\eta_l^2 L^2) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2 \sigma_i^2}{b_i} + 12\eta_l^2 \zeta_i^2
\end{aligned}$$

Bounding local uniform Stability: Now, we can combine these two cases. Our bound relies on the probability of whether the perturbed samples are involved. Thus, we have:

$$\begin{aligned}
&\mathbb{E}_{\mathcal{D}_i, \mathcal{D}'_i} \|\mathbf{x}_i^{k+1} - \tilde{\mathbf{x}}_i^{k+1}\|^2 \\
&= P(\tilde{z} \notin \xi) \cdot \mathbb{E}[\|\mathbf{x}_i^{k+1} - \tilde{\mathbf{x}}_i^{k+1}\|^2 | \tilde{z} \notin \xi] + P(\tilde{z} \in \xi) \cdot \mathbb{E}[\|\mathbf{x}_i^{k+1} - \tilde{\mathbf{x}}_i^{k+1}\|^2 | \tilde{z} \in \xi] \\
&\leq (1 - \frac{b_i}{n_i}) \left((1 + 2\eta_l L + 2\eta_l^2 L^2) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + \frac{4\eta_l^2 \sigma_i^2}{b_i} \right) \\
&\quad + \frac{b_i}{n_i} \left((1 + 2\eta_l L + 6\eta_l^2 L^2) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + 4\eta_l^2 \left(\frac{\sigma_i^2}{b_i} + 3\zeta_i^2 \right) \right) \\
&= (1 + 2\eta_l L + 2(1 + 2b_i/n_i)\eta_l^2 L^2) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + 4 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i \zeta_i^2}{n_i} \right) \eta_l^2 \quad \triangleright b_i/n_i \leq 1/2 \\
&\leq (1 + 2\eta_l L + 4\eta_l^2 L^2) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + 4 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i \zeta_i^2}{n_i} \right) \eta_l^2 \quad \triangleright 1 + 2\eta_l L \leq 2 \\
&\leq (1 + 4\eta_l L) \mathbb{E} \|\mathbf{x}_i^k - \tilde{\mathbf{x}}_i^k\|^2 + 4 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i \zeta_i^2}{n_i} \right) \eta_l^2,
\end{aligned}$$

where ξ denotes the mini-batch data sampled from the local dataset. According to the model merge rule, the models fine-tuned for different tasks are all trained from the same pre-trained model, i.e., $\mathbf{x}_i^0 = \mathbf{x}^0$. Then, we unroll the above equation over local steps from $K_i - 1$ down to $k = 0$ ($K_i > 1$):

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}_i, \mathcal{D}'_i} \|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 &\leq (1 + 4\eta_l L)^{K_i} \mathbb{E} \|\mathbf{x}^0 - \tilde{\mathbf{x}}^0\|^2 + \frac{(1 + 4\eta_l L)^{K_i} - 1}{4\eta_l L} \cdot 4\eta_l^2 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i \zeta_i^2}{n_i} \right) \\
&= \frac{(1 + 4\eta_l L)^{K_i} - 1}{4\eta_l L} \cdot 4\eta_l^2 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i \zeta_i^2}{n_i} \right)
\end{aligned}$$

Here we use $\mathbf{x}^0 = \tilde{\mathbf{x}}^0$, which holds because both the perturbed and the unperturbed parameters are fine-tuned from the same pre-trained model. Assume the learning rate satisfies $\frac{1}{8KL} \geq \eta_l \geq \frac{1}{16KL}$, then we have $(1 + 4\eta_l L)^{K_i} \leq 2$ and $\frac{1}{\eta_l L} \leq 16K_i$. We get:

$$\mathbb{E}_{\mathcal{D}_i, \mathcal{D}'_i} \|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 \leq 16K_i \eta_l^2 \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i \zeta_i^2}{n_i} \right)$$

Then we complete the proof.

Step 2: Proofs of Global Stability

Since model merging involves only a single aggregation step, different merging methods such as Task Arithmetic [31] and AdMerging [93] essentially learn different merging coefficients. Since higher-order model merging methods such as AdaMerging are capable of automatically capturing the differences between tasks, they can adaptively adjust the merging coefficients to achieve better performance. This is equivalent to modifying the objective function from the standard form $f(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{x})$ to a weighted form $f(\mathbf{x}) = \sum_{i=1}^N \lambda_i f_i(\mathbf{x})$. We will focus on discussing the objective function discrepancies caused by heterogeneity in [step 3](#). The key difference among various model merging approaches lies in how they learn the parameters λ_i . To generalize our theoretical results, we will adopt the more general loss function $f(\mathbf{x}) = \sum_{i=1}^N \lambda_i f_i(\mathbf{x})$ in the analysis of Global Stability, where $\sum_{i=1}^N \lambda_i = 1$ according to [Assumption 4](#).

Specifically, the model merge step can be formally expressed as $\mathbf{x}_{avg} = \sum_{i=1}^N \lambda_i \mathbf{x}_i^{K_i}$, $\tilde{\mathbf{x}}_{avg} = \sum_{i=1}^N \lambda_i \tilde{\mathbf{x}}_i^{K_i}$. Then, the Global Stability can be represented as $\mathbb{E}_{\mathcal{D}, \mathcal{D}^{(j)}} \|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2$. Next, we derive an upper bound for the Global Stability.

$$\begin{aligned}
& \mathbb{E}_{\mathcal{D}, \mathcal{D}^{(j)}} \|\mathbf{x}_{avg} - \tilde{\mathbf{x}}_{avg}\|^2 \\
&= \mathbb{E}_{\mathcal{D}, \mathcal{D}^{(j)}} \left\| \sum_{i=1}^N \lambda_i (\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}) \right\|^2 \quad \triangleright \text{Jensen's Inequality} \\
&\leq \sum_{i=1}^N \lambda_i \mathbb{E}_{\mathcal{D}, \mathcal{D}^{(j)}} \|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 \\
&= \sum_{i=1}^N \lambda_i \left(\mathcal{P}(\tilde{z} \in \mathcal{D}_i) \cdot \mathbb{E} \left[\|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 | \tilde{z} \in \mathcal{D}_i \right] + \mathcal{P}(\tilde{z} \notin \mathcal{D}_i) \cdot \mathbb{E} \left[\|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 | \tilde{z} \notin \mathcal{D}_i \right] \right) \\
&= \sum_{i=1}^N \lambda_i \left(\underbrace{\frac{n_i}{n} \cdot \mathbb{E} \left[\|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 | \tilde{z} \in \mathcal{D}_i \right]}_{\text{Lemma 3}} + \left(1 - \frac{n_i}{n}\right) \cdot \underbrace{\mathbb{E} \left[\|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 | \tilde{z} \notin \mathcal{D}_i \right]}_{\text{Given Below}} \right)
\end{aligned}$$

Then, we let $b_i = 0$ for (8) denote the local stability without the perturbed sample (i.e., $\tilde{z} \notin \mathcal{D}_i$). In this case, the local mini-batch SGD updates are amplified due to local stochastic gradient variance and cumulative SGD steps as:

$$\mathbb{E} \left[\|\mathbf{x}_i^{K_i} - \tilde{\mathbf{x}}_i^{K_i}\|^2 | \tilde{z} \notin \mathcal{D}_i \right] \leq 16K_i \eta_l^2 \frac{\sigma_i^2}{n_i}$$

Combining the above equations, we have

$$\mathbb{E}_{\mathcal{D}, \mathcal{D}^{(j)}} \|\mathbf{x}_{avg} - \tilde{\mathbf{x}}_{avg}\|^2 \leq 16\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2 \right) \quad (10)$$

This concludes the proof of global stability.

Step 3: Convergence Analysis (Optimization Error)

Due to the heterogeneity in training resources or algorithmic hyperparameters during the fine-tuning stage across different tasks (e.g., different tasks may use different learning rates η_l , epochs K_i , optimizers, or batch sizes b_i), there may exist significant discrepancies among the parameters learned by each task. Whether merging such diverse model parameters can still lead to convergence is an important question worth exploring. To address this, we follow the analytical framework proposed by Wang et al. [75], with necessary modifications to suit the context of model merging. We will first introduce some notation and then analyze the theoretical implications of such heterogeneity.

Notations: First, we define the following matrix of stochastic gradients for each task i :

$$\mathbf{G}_i = [\mathbf{g}_i(\mathbf{x}_i^0), \mathbf{g}_i(\mathbf{x}_i^1), \dots, \mathbf{g}_i(\mathbf{x}_i^{K_i-1})] \in \mathbb{R}^{d \times K_i}$$

Here, \mathbf{g}_i denotes a stochastic gradient of $f_i(\mathbf{x})$. Next, we define the normalized learning rate vector.

$$\mathbf{a}_i = \left[\frac{\eta_i^0}{\eta_l}, \frac{\eta_i^1}{\eta_l}, \dots, \frac{\eta_i^{K_i-1}}{\eta_l} \right]^T \in \mathbb{R}^{K_i}$$

Here, η_l denotes a constant learning rate, while η_i^t , with $t \in [K_i]$, $i \in [N]$, can represent learning rates under any strategy,

such as exponential decay.

$$\begin{aligned}
\mathbf{x}_{avg} &= \mathbf{x}_0 + \frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i^{K_i} - \mathbf{x}_0) \\
&= \mathbf{x}_0 - \frac{1}{N} \sum_{i=1}^N \eta_l \mathbf{G}_i \mathbf{a}_i \\
&= \mathbf{x}_0 - \underbrace{\left(\frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i\|_1 \right)}_{\tau_{\text{eff}}: \text{effective local steps}} \sum_{i=1}^N \eta_l \underbrace{\left(\frac{\frac{1}{N} \|\mathbf{a}_i\|_1}{\frac{1}{N} \sum_{i=1}^N \|\mathbf{a}_i\|_1} \right)}_{\lambda_i: \text{weight}} \underbrace{\left(\frac{\mathbf{G}_i \mathbf{a}_i}{\|\mathbf{a}_i\|_1} \right)}_{\mathbf{d}_i: \text{normalized gradient}}
\end{aligned} \tag{11}$$

Impact of Heterogeneity: Unlike homogeneous training—where each task uses the same learning rate, number of epochs, batch size, etc.—heterogeneous training results in parameter updates that differ across tasks. According to the mathematical implication of Equation (11), such heterogeneity effectively transforms the original optimization objective $f(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{x})$ into a surrogate objective $\tilde{f}(\mathbf{x}) = \sum_{i=1}^N \lambda_i f_i(\mathbf{x})$, where the coefficients λ_i are defined as in Equation (11). Therefore, training heterogeneity introduces inconsistency in the objective function, making the theoretical analysis of model merging under heterogeneous conditions more complex and also more meaningful. In the following analysis, we first present a more general theorem, and then provide a simplified version that establishes an upper bound on the optimization error under the assumption of an SGD optimizer. In the following analysis, we first present a more general theorem, which is adapted from the work of [75]. However, we have made several modifications to their original result, as our Assumption 3 differs from Assumption 3 in [75]. Nevertheless, we show that our assumption implies theirs, and the proof is provided as follows:

Assumption 5 (Bounded Dissimilarity in [75]) *Let the merging coefficients satisfy Assumption 4, and suppose there exist constants $\beta^2 \leq 1$ and $\kappa^2 \leq 1$ such that*

$$\sum_{i=1}^N \lambda_i \|\nabla f_i(\mathbf{x})\|^2 \leq \beta^2 \left\| \sum_{i=1}^N \lambda_i \nabla f_i(\mathbf{x}) \right\|^2 + \kappa^2.$$

Lemma 4 *Assumption 3 in this paper implies Assumption 5 in the work of Wang et al. [75].*

Proof 3 *As long as we can identify the parameters β^2 and κ^2 in Assumption 5, we establish the above lemma. We have*

$$\begin{aligned}
\sum_{i=1}^N \lambda_i \|\nabla f_i(\mathbf{x})\|^2 &= \sum_{i=1}^N \lambda_i \|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x}) + \nabla f(\mathbf{x})\|^2 \\
&\leq 2 \sum_{i=1}^N \lambda_i \|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 + 2 \sum_{i=1}^N \lambda_i \|\nabla f(\mathbf{x})\|^2 \quad \triangleright \text{Assumption 4 \& 3} \\
&\leq 2 \left\| \sum_{i=1}^N \lambda_i \nabla f_i(\mathbf{x}) \right\|^2 + 2 \sum_{i=1}^N \lambda_i \zeta_i^2
\end{aligned}$$

By setting $\beta^2 = 2$ and $\kappa^2 = 2 \sum_{i=1}^N \lambda_i \zeta_i^2$, the proof is complete.

Theorem 4 (Convergence to the Surrogate Objective $\tilde{f}(\mathbf{x})$ Stationary Point from [75]) *Under Assumptions 1, 2, and 5, let $\bar{K} = \frac{1}{N} \sum_{i=1}^N K_i$ and $\eta_l = \sqrt{\frac{N}{\bar{K}}}$. Then, the gradient of the surrogate function $\tilde{f}(\mathbf{x}) = \sum_{i=1}^N \lambda_i f_i(\mathbf{x})$ at the averaged model \mathbf{x}_{avg} is bounded as follows:*

$$\mathbb{E}[\|\nabla \tilde{f}(\mathbf{x}_{avg})\|^2] \leq \frac{4(\tilde{f}(\mathbf{x}_0) - \tilde{f}_{\text{inf}})(\bar{K}/\tau_{\text{eff}})}{\sqrt{N\bar{K}}} + \frac{4L\sigma^2 A_1}{\sqrt{N\bar{K}}} + \frac{6NL^2\sigma^2 A_2}{\bar{K}} + \frac{12NL^2\zeta^2 A_3}{\bar{K}}. \tag{12}$$

Where $\tilde{f}_{\text{inf}} = \min_{\mathbf{x}} \tilde{f}(\mathbf{x})$, $\sigma = \max_i \sigma_i$, $A_1 = \tau_{\text{eff}} N \sum_{i=1}^N \frac{\lambda_i^2 \|\mathbf{a}_i\|_2^2}{\|\mathbf{a}_i\|_1^2}$, $A_2 = \sum_{i=1}^N \lambda_i (\|\mathbf{a}_i\|_2^2 - a_{i,-1}^2)$ and $A_3 = \max_i \{\|\mathbf{a}_i\|_1 (\|\mathbf{a}_i\|_1 - a_{i,-1})\}$, where $a_{i,-1}$ denotes the last coordinate of the vector \mathbf{a}_i . Furthermore, let the right-hand

side of Equation 12 be denoted as ϵ . Then, we have the following inequality, which ensures that even under heterogeneous conditions, the original objective function $f(\mathbf{x})$ can still converge, although its bound is not as tight as that in Equation 12:

$$\mathbb{E}[\|\nabla f(\mathbf{x}_{avg})\|^2] \leq 2[\chi_{\mathbf{p}|\boldsymbol{\lambda}}^2 + 1]\epsilon + 2\chi_{\mathbf{p}|\boldsymbol{\lambda}}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 \quad (13)$$

where $\chi_{\mathbf{p}|\boldsymbol{\lambda}}^2 = \sum_{i=1}^N \frac{(\frac{1}{N} - \lambda_i)^2}{\lambda_i^2}$ is the chi-square divergence between the weight coefficient vectors $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_N] \in \mathbb{R}^N$ and $\mathbf{p} = [\frac{1}{N}, \frac{1}{N}, \dots, \frac{1}{N}] \in \mathbb{R}^N$.

In fact, different choices of optimizers in Theorem 4 yield different forms of the weighting vector \mathbf{a}_i . In the work of Wang et al. [75], several illustrative examples are provided. For instance, when the optimizer incorporates a proximal term, the corresponding \mathbf{a}_i is given by

$$\mathbf{a}_i = [(1 - \alpha)^{K_i - 1}, (1 - \alpha)^{K_i - 2}, \dots, (1 - \alpha), 1] \in \mathbb{R}^{K_i},$$

where $\alpha = \alpha\mu$ and μ is a tunable parameter. In this case, the effective number of iterations and the associated weights are defined as

$$\tau_{\text{eff}} = \sum_{i=1}^N \frac{1 - (1 - \alpha)^{K_i}}{\alpha N}, \quad \lambda_i = \frac{1 - (1 - \alpha)^{K_i} / N}{\sum_{i=1}^N [1 - (1 - \alpha)^{K_i}] / N}.$$

Moreover, Wang et al. [75] also derive explicit forms of \mathbf{a}_i , τ_{eff} , and λ_i when using other optimizers such as SGD with learning rate decay or SGD with momentum.

In summary, Theorem 4 provides a general analytical framework accommodating a variety of optimization strategies. For the sake of simplicity and to align with the hyperparameter settings used in the *Step 2 Global Stability Analysis*, we proceed with a simplified version assuming that each task is optimized using standard SGD. Under this setting, we have

$$\mathbf{a}_i = [1, 1, \dots, 1] \in \mathbb{R}^{K_i}, \quad \tau_{\text{eff}} = \bar{K} = \frac{1}{N} \sum_{i=1}^N K_i, \quad \|\mathbf{a}_i\|_1 = K_i, \quad \lambda_i = \frac{K_i}{N\bar{K}},$$

which leads directly to the following corollary of Theorem 4.

Corollary 1 (Adopt the same SGD optimizer for all tasks from Theorem 4) Under Assumptions 1, 2, and 5, let $\bar{K} = \frac{1}{N} \sum_{i=1}^N K_i$, $\eta = \sqrt{\frac{N}{\bar{K}}}$ and all task use the SGD optimizer. Then, the gradient of the surrogate function $\tilde{f}(\mathbf{x}) = \sum_{i=1}^N \lambda_i f_i(\mathbf{x})$ at the averaged model \mathbf{x}_{avg} is bounded as follows:

$$\mathbb{E}[\|\nabla \tilde{f}(\mathbf{x}_{avg})\|^2] \leq \frac{4(\tilde{f}(\mathbf{x}_0) - \tilde{f}_{\text{inf}})}{\sqrt{N\bar{K}}} + \frac{4L\sigma^2 A_1}{\sqrt{N\bar{K}}} + \frac{6NL^2\sigma^2 A_2}{\bar{K}} + \frac{12NL^2\zeta^2 A_3}{\bar{K}}. \quad (14)$$

Where $\tilde{f}_{\text{inf}} = \min_{\mathbf{x}} \tilde{f}(\mathbf{x})$, $\sigma = \max_i \{\sigma_i\}$, $\zeta = \max_i \{\zeta_i\}$, $A_1 = N \sum_{i=1}^N \frac{\bar{K}}{K_i} \lambda_i^2$, $A_2 = \sum_{i=1}^N \lambda_i (K_i - 1)$ and $A_3 = \max_i \{K_i (K_i - 1)\}$. Furthermore, let the right-hand side of Equation 14 be denoted as ϵ_{sgd} . Then, we have the following inequality, which ensures that even under heterogeneous conditions, the original objective function $f(\mathbf{x})$ can still converge, although its bound is not as tight as that in Equation 14:

$$\mathbb{E}[\|\nabla f(\mathbf{x}_{avg})\|^2] \leq 2[\chi_{\mathbf{p}|\boldsymbol{\lambda}}^2 + 1]\epsilon_{sgd} + 2\chi_{\mathbf{p}|\boldsymbol{\lambda}}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 \quad (15)$$

where $\chi_{\mathbf{p}|\boldsymbol{\lambda}}^2 = \sum_{i=1}^N \frac{(\frac{1}{N} - \lambda_i)^2}{\lambda_i^2}$ is the chi-square divergence between the weight coefficient vectors $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \dots, \lambda_N] \in \mathbb{R}^N$ and $\mathbf{p} = [\frac{1}{N}, \frac{1}{N}, \dots, \frac{1}{N}] \in \mathbb{R}^N$.

Step 4: Excess Risk Analysis via Joint Minimization

According to the definition of excess error in Equation 3, we have $\mathcal{E}(\mathbf{x}) = \mathcal{E}_G(\mathbf{x}) + \mathcal{E}_O(\mathbf{x})$. Furthermore, we assume that the optimization error satisfies $\mathcal{E}_O(\mathbf{x}_{avg}) \leq C \cdot \mathbb{E}[\|\nabla f(\mathbf{x}_{avg})\|^2]$. When $C = \frac{1}{2\mu}$, this assumption degenerates to the Polyak-Łojasiewicz (PL) condition proposed in [33]. However, for non-convex optimization analysis, the PL condition is sometimes

considered too strong. A milder assumption is the Kurdyka-Łojasiewicz (KL) condition, which yields $\mathcal{E}_O(\mathbf{x}_{avg}) \leq C^{\frac{1}{2\theta}} \cdot \mathbb{E}\|\nabla f(\mathbf{x}_{avg})\|^{\frac{1}{\theta}}$, where f is assumed to be smooth and $\theta \in [0, 1)$. Notably, when $\theta = \frac{1}{2}$, the KL condition recovers the PL condition. To maintain generality, we do not explore how the KL parameter θ improves convergence [37] or generalization [7], thereby avoiding strong assumptions.

Combining Lemma 2, Equation 3, and the assumption $\mathcal{E}_O(\mathbf{x}_{avg}) \leq C \cdot \mathbb{E}\|\nabla f(\mathbf{x}_{avg})\|^2$, we derive the following stability-based expression for the excess error:

$$\mathcal{E}(\mathbf{x}_{avg}) \leq \frac{L + \gamma}{2} \underbrace{\mathbb{E}\|\mathbf{x}_{avg} - \tilde{\mathbf{x}}_{avg}\|^2}_{\text{Model stability}} + \left(\frac{1}{2\gamma} + C\right) \underbrace{\mathbb{E}\|\nabla f(\mathbf{x}_{avg})\|^2}_{\text{Gradient norm}}.$$

Substituting inequality 10 and 14 into the above equation, we obtain:

$$\mathcal{E}(\mathbf{x}_{avg}) \leq 8(L + \gamma)\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2\right) + \left(\frac{1}{\gamma} + 2C\right) \left[\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}|\lambda}^2 + 1)\epsilon_{sgd} \right] \quad (16)$$

Where ϵ_{sgd} is defined in inequality (15).

Next, we will determine the optimal γ for the above equation to minimize the upper bound on the right-hand side. After a straightforward differentiation, we obtain that when γ takes the following form,

$$\gamma^* = \sqrt{\frac{\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}|\lambda}^2 + 1)\epsilon_{sgd}}{8\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2\right)}}$$

Consequently, the tightest bound of $\mathcal{E}(\mathbf{x}_{avg})$ with respect to γ is:

$$\begin{aligned} \mathcal{E}(\mathbf{x}_{avg}) &\leq 8L\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2\right) + 2C \left[\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}|\lambda}^2 + 1)\epsilon_{sgd} \right] \\ &\quad + 4\sqrt{2}\eta_l \sqrt{\sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2\right) \cdot \left[\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}|\lambda}^2 + 1)\epsilon_{sgd} \right]} \end{aligned}$$

To facilitate the analysis, we apply the basic inequality $2\sqrt{ab} \leq a + b$ to the radical term in the above expression, yielding:

$$\begin{aligned} \mathcal{E}(\mathbf{x}_{avg}) &\leq 8(L + 1)\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2\right) \\ &\quad + (2C + 1) \left[\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}|\lambda}^2 + 1)\epsilon_{sgd} \right] \end{aligned} \quad (17)$$

A.1.1. Generalization Analysis of Classical Model Merge Methods

The excess risk bound derived in Equation (17) provides a powerful theoretical lens through which we can analyze and understand the effectiveness of various model merging strategies. The bound reveals a fundamental trade-off: minimizing the risk $\mathcal{E}(\mathbf{x}_{avg})$ requires jointly controlling the model stability and the optimization error, both of which are intricately linked to the merging coefficients λ and the inherent heterogeneity of the tasks, quantified by ζ_i^2 and the chi-square divergence $\chi_{\mathbf{p}|\lambda}^2$.

Different model merging methods can be interpreted as different strategies for selecting the coefficients λ_i , thereby navigating this trade-off in distinct ways.

1. Simple Averaging and Task Arithmetic [31]:

- **Strategy:** The most straightforward approach is to perform a uniform average of model parameters, which corresponds to setting $\lambda_i = 1/N$ for all i . This is the default strategy in many foundational works. Task Arithmetic [31] operates on task vectors ($\tau_i = \mathbf{x}_i - \mathbf{x}_0$) but often uses simple scalar weights, which are mathematically equivalent to a weighted average of model parameters.

- **Theoretical Interpretation:** The primary advantage of uniform averaging ($\lambda_i = 1/N$) is that it perfectly aligns the surrogate objective with the original objective. In this case, the weight vector λ is identical to the uniform distribution vector p , causing the chi-square divergence term $\chi_{p||\lambda}^2$ to become zero. This completely eliminates the penalty associated with optimizing a surrogate objective, as seen in the second major term of Equation (17). However, this strategy is naive to task heterogeneity. It assigns equal importance to all tasks, regardless of their dissimilarity. If one task is a significant outlier (i.e., has a very large ζ_i^2), its negative impact is not mitigated, potentially inflating both the stability term and the optimization error term ($\sum \lambda_i \zeta_i^2$).
2. **Sparsity-Based Merging (DARE & TIES-Merging):**
 - **Strategy:** Methods like DARE (Drop and Rescale) [95] and TIES-Merging [90] introduce sparsity into the task vectors before merging. TIES-Merging identifies and retains only the most significant parameter changes, resetting the rest. DARE randomly drops a fraction of the parameter changes from each task vector.
 - **Theoretical Interpretation:** These methods address the problem of "parameter interference," where different tasks pull the same parameter in conflicting directions. In our framework, this interference is a major contributor to the heterogeneity term ζ_i^2 . By enforcing sparsity, DARE and TIES-Merging effectively create modified task vectors τ'_i . This process can be seen as a form of projection that filters out conflicting or less important updates. Consequently, the effective heterogeneity among the sparse vectors is reduced, leading to a smaller ζ_i^2 in the final merged model. This directly reduces both the stability and optimization error terms in our bound, thereby improving generalization. They mitigate task conflicts at a fine-grained, parameter level rather than at a coarse, task level.
 3. **Adaptive Merging (AdaMerging):**
 - **Strategy:** Higher-order methods like AdaMerging [93] learn the merging coefficients λ_i by solving an auxiliary optimization problem on a small, held-out dataset. The goal is to find weights that produce a merged model with the lowest loss on this proxy dataset.
 - **Theoretical Interpretation:** AdaMerging provides a principled approach to actively minimize the generalization bound. Its strategy directly confronts the trade-off highlighted by our theory.
 - **Tackling Heterogeneity:** The auxiliary optimization process is designed to identify and down-weight outlier or conflicting tasks. If a task j is highly dissimilar to others (large ζ_j^2), including it with a large weight λ_j would likely increase the loss on the proxy dataset. Therefore, the optimization will naturally assign a smaller λ_j to it. This explicitly minimizes the impact of the heterogeneity-dependent term $\sum_{i=1}^N \lambda_i \zeta_i^2$.
 - **Navigating the Trade-off:** By learning non-uniform weights, AdaMerging accepts a non-zero surrogate objective penalty ($\chi_{p||\lambda}^2 > 0$). However, it does so only when the benefit of reducing the impact of heterogeneity ($\sum \lambda_i \zeta_i^2$) outweighs the cost. It finds an empirically optimal balance between aligning with the uniform objective and adapting to the specific geometry of the task vectors. In essence, AdaMerging uses a data-driven approach to find a λ that minimizes a proxy for the entire bound in Equation 17, rather than greedily minimizing just one of its components.
 4. **Pre-Merging Methods (Weight Alignment):**
 - **Strategy:** Methods like Git Re-basin [1] operate before the merge. They find permutation symmetries between the neurons of different models and permute the weights of one model to geometrically align with another, without changing its functional output. The goal is to place models in the same loss basin before averaging.
 - **Theoretical Interpretation:** From our framework's perspective, weight alignment is a proactive strategy to *reduce the initial task heterogeneity* ζ_i^2 . By aligning models geometrically, these methods reduce the dissimilarity between their parameter vectors. This ensures that the subsequent merging step begins with a set of models that are already more concordant, effectively tightening the starting conditions for our excess error bound and leading to a more stable and better-performing final model.
 5. **Post-Merging Methods (Representation Surgery):**
 - **Strategy:** These methods, exemplified by Representation Surgery [92], operate on the already-merged model x_{avg} . They introduce a lightweight, task-specific module that calibrates the model's internal representations rather than altering the core merged parameters.
 - **Theoretical Interpretation:** Post-merging methods offer a sophisticated strategy for managing the compromises inherent in our bound. They implicitly treat the solution as a decoupled system.
 - **Core Model:** The main merging process finds a robust, "common ground" model x_{avg} that minimizes the bulk of the excess error, primarily the components shared across tasks.
 - **Corrective Module:** The post-hoc module is then tasked with handling the *residual excess error*, specifically the portion driven by high task heterogeneity (ζ_i^2) that could not be fully resolved in the shared parameter space without causing interference. By compensating for this error in the representation space, it provides task-specific calibration

without destabilizing the carefully balanced merged model.

In summary, our theoretical framework successfully explains the hierarchy of model merging methods. While simple averaging minimizes the surrogate objective penalty at the risk of being susceptible to outliers, more advanced methods like DARE and AdaMerging employ sophisticated strategies—either by reducing parameter-level interference or by adaptively learning task-level coefficients—to actively mitigate the negative effects of task heterogeneity, ultimately leading to a tighter excess risk bound and better generalization.

A.1.2. Deeper Analysis of the Generalization Bound and Hyperparameters

While understanding how different merging methods correspond to choices of λ is crucial, a deeper analysis of the hyperparameters within the excess risk bound in Equation 17 reveals fundamental trade-offs in the fine-tune-then-merge paradigm. We now dissect the role of key hyperparameters.

The final bound is driven by two main components: a *stability term*, scaled by $8(L+1)\eta_l^2$, and an *optimization error term*, scaled by $(2C+1)$.

$$\mathcal{E}(\mathbf{x}_{avg}) \leq \underbrace{8(L+1)\eta_l^2 \sum_{i=1}^N \lambda_i K_i \left(\frac{\sigma_i^2}{n_i} + \frac{3b_i}{n_i} \zeta_i^2 \right)}_{\text{Stability Component}} + \underbrace{(2C+1) \left[\chi_{\mathbf{p}|\lambda}^2 \sum_{i=1}^N \lambda_i \zeta_i^2 + (\chi_{\mathbf{p}|\lambda}^2 + 1) \epsilon_{sgd} \right]}_{\text{Optimization Error Component}}$$

Let's analyze the impact of each hyperparameter:

- Number of Local Fine-Tuning Steps (K_i):** The number of local epochs or steps is a critical parameter that embodies the core trade-off between optimization and generalization.
 - Mathematical Impact:** K_i appears linearly in the numerator of the stability component, suggesting that more steps can worsen stability. Conversely, the optimization error term ϵ_{sgd} generally decreases with the average number of steps, \bar{K} , typically at a rate of $O(1/\sqrt{\bar{K}})$ or $O(1/\bar{K})$.
 - Interpretation (The Optimization-Generalization Trade-off):**
 - A **small** K_i keeps the stability term low. The model stays close to the pre-trained initialization (\mathbf{x}_0), and its final state is less sensitive to individual data points in its local dataset \mathcal{D}_i . However, a small K_i leads to a large ϵ_{sgd} , meaning the model is poorly optimized (under-fitting).
 - A **large** K_i reduces the optimization error ϵ_{sgd} , pushing the model closer to the minimum of its local objective. However, this comes at the cost of increased instability. With more steps, the model has more capacity to over-specialize on \mathcal{D}_i , causing a small data perturbation to lead to a larger divergence in the final parameters $\mathbf{x}_i^{K_i}$, which harms generalization.
 - Practical Insight:** This trade-off implies that an optimal number of fine-tuning steps exists. Depending on which factor—optimization error or instability—dominates, performance may first increase and then decrease (an inverted U-shape), or it could decrease monotonically if the stability cost is immediately too high. Overtraining leads to over-specialized models that are difficult to merge, a phenomenon our bound quantifies as a loss of algorithmic stability. Therefore, we recommend incorporating more metrics to guide the selection of the optimal K_i in practical tasks. In general, it is advisable to employ multiple monitoring mechanisms to determine whether the optimization process has converged. Once convergence is detected, training should be stopped immediately, as continuing to train will increase the stability term, while the reduction in optimization error becomes marginal and outweighed by the degradation in stability.
- Batch Size (b_i):** The batch size exhibits a subtle, dual role whose practical effect is conditioned by other hyperparameters.
 - Mathematical Impact:** b_i appears in the numerator of the heterogeneity-driven part of the stability term ($3b_i\zeta_i^2/n_i$), suggesting a larger batch size can harm stability. However, the entire stability term is scaled by η_l^2 .
 - Interpretation (A Conditional Trade-off):**
 - In principle, b_i presents a trade-off. A **small** b_i keeps the term $3b_i\zeta_i^2/n_i$ small but can lead to noisy gradients that slow convergence, affecting the optimization error ϵ_{sgd} .
 - A **large** b_i reduces the variance of stochastic gradients, typically improving the optimization error ϵ_{sgd} . However, it linearly increases the stability cost term $3b_i\zeta_i^2/n_i$, making the algorithm more sensitive to data perturbations.
 - The **decisive factor** is the learning rate η_l . Since the stability cost is scaled by η_l^2 , a small learning rate can heavily suppress this negative term, making the positive effect on optimization the dominant factor.
 - Practical Insight:** Our bound provides a theoretical rationale for why the observed effect of batch size depends on the experimental setting. While a trade-off exists in theory, for fine-tuning regimes with small learning rates, the benefits

of improved gradient estimation from larger batch sizes are expected to dominate, leading to monotonically better performance. Therefore, we recommend using a large batch size for fine-tuning in practical tasks whenever possible.

3. **Dataset Size (n_i):** The amount of data per task has a clear and unambiguously beneficial effect.
 - **Mathematical Impact:** n_i appears in the denominator of the entire stability term, scaling both the variance component (σ_i^2/n_i) and the heterogeneity component ($3b_i\zeta_i^2/n_i$).
 - **Interpretation (A "Free Lunch" for Generalization):** Unlike other hyperparameters that present a trade-off, increasing n_i strictly tightens the generalization bound by enhancing model stability without adversely affecting the optimization error component. The influence of any single data point is diminished as the dataset grows, reducing the model's sensitivity to data perturbations.
 - **Practical Insight:** This aligns with the foundational principle of statistical learning: more data leads to better generalization. Our bound confirms this holds true in the context of model merging, predicting a direct and monotonic improvement in performance with more data. Therefore, we recommend using as much data as possible for fine-tuning in practical tasks.
4. **Learning Rate (η_l):** The learning rate is a dominant factor in controlling stability.
 - **Mathematical Impact:** η_l appears as a squared term, η_l^2 , multiplying the entire stability component.
 - **Interpretation (The Convergence vs. Stability Trade-off):**
 - A **small** η_l is essential for good generalization. The η_l^2 term indicates that the stability bound is highly sensitive to the learning rate. A smaller step size ensures that any perturbation in the gradient results in only a small change in the final parameters, thus promoting stability.
 - A **large** η_l can be catastrophic for generalization. It can cause the stability component to explode, leading to an unstable algorithm whose output varies wildly with small data changes, potentially even causing model collapse after merging.
 - **Practical Insight:** The bound provides strong theoretical justification for using small learning rates during fine-tuning. This is crucial not only to avoid catastrophic forgetting but also, as our theory shows, to ensure the resulting model is stable enough to be effectively merged and to generalize well. Therefore, we recommend using a small learning rate for fine-tuning in practical tasks.
5. **Number of Tasks (N):** The number of tasks being merged introduces a trade-off between a minor optimization benefit and a significant heterogeneity cost.
 - **Mathematical Impact:** N appears in the denominator of the optimization error term ϵ_{sgd} (improving at a slow rate of $O(1/\sqrt{N})$), but the total task heterogeneity, which impacts several terms, grows linearly with N .
 - **Interpretation (The Optimization vs. Heterogeneity Imbalance):**
 - A **larger** N provides a small, slowly-improving benefit to the optimization error. However, this benefit's leading coefficient, related to the initial suboptimality gap, is small for models finetuned from strong pretrained checkpoints.
 - Conversely, as more tasks are added, the average task heterogeneity increases. This "parameter pollution" from conflicting tasks creates a significant penalty that accumulates rapidly. There is a fundamental imbalance: a minor benefit with a small coefficient is pitted against a substantial penalty that grows quickly.
 - **Practical Insight:** Our bound predicts that for a diverse set of tasks, the performance degradation from escalating task heterogeneity will overwhelm the marginal gains in optimization. This leads to the expectation that generalization performance will decrease as N increases, highlighting a key limitation of simple averaging in many-task scenarios. Therefore, we recommend minimizing the inclusion of models unrelated to the target task during merging, as reducing the number of merged models can improve overall performance.

B. Experiment Details

B.1. Datasets Setup

We evaluate our methods on a diverse set of 20 downstream tasks spanning multiple domains. These include natural scene classification (SUN397 [87]), fine-grained object recognition (Stanford-Cars [40], Oxford Flowers102 [56]), remote sensing (RESISC45 [9], EuroSAT [28]), digit and character recognition (MNIST [42], SVHN [55], EMNIST [15], KMNIST [13]), traffic sign recognition (GTSRB [64]), texture classification (DTD [12]), medical imaging (PCAM [72]), facial expression recognition (FER2013 [24]), pet breed classification (Oxford-IIIT-Pet [59]), general object recognition (CIFAR10, CIFAR100 [41], STL10 [14], Food101 [4]), fashion product classification (Fashion-MNIST [86]), and sentiment analysis (Rendered-SST2 [61, 63]). This selection covers a wide range of visual and textual tasks, ensuring comprehensive evaluation of model merging strategies under heterogeneous domains and data distributions.

C. More Results

C.1. The Average Performance of Individual Models Before Merging

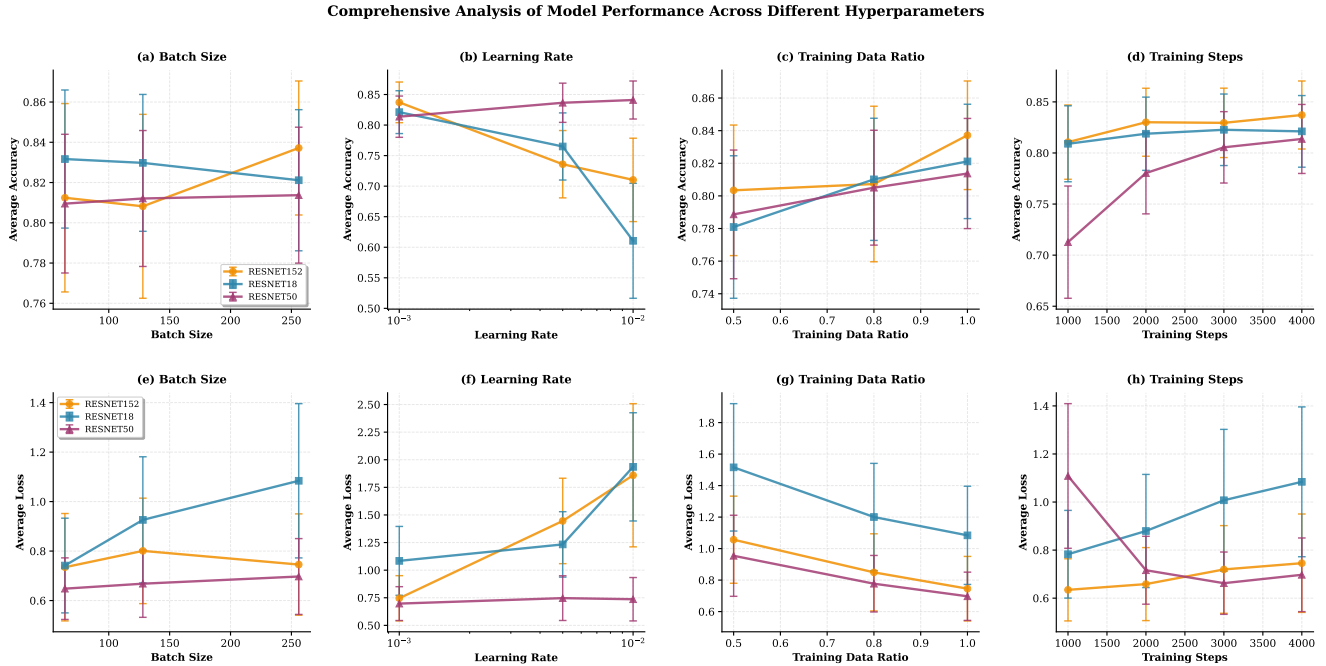


Figure 3. **Hyperparameter Impact on Model Merging Performance.** Comprehensive analysis of four key finetuning hyperparameters across three ResNet architectures. Top row (a-d): accuracy vs. batch size, learning rate, training data ratio, and training steps. Bottom row (e-h): corresponding loss curves. Each experiment varies one hyperparameter while fixing others at defaults (batch size=256, lr=0.001, data ratio=1.0, steps=4000). Results averaged over 20 vision datasets with error bars showing standard error.

C.2. Single-Task Finetuning Performance

To provide a comprehensive baseline for our merging experiments, the following table details the performance of each backbone model after being individually fine-tuned on its respective task. These accuracy scores represent the "expert" model performance before any model merging techniques are applied.

Table 2. **Single-Task Finetuning Accuracy.** Performance (top-1 accuracy) of ResNet-18, ResNet-50, and ResNet-152 models after being individually fine-tuned on each of the 20 downstream vision tasks. **All models were fine-tuned for 4000 steps with a learning rate of 0.001, a batch size of 256, and using the full training dataset (data ratio of 1.0).** The highest accuracy for each task is highlighted in **bold**. This data serves as the baseline for the expert models used in our merging experiments.

Task	ResNet-18	ResNet-50	ResNet-152
cifar10	0.9404	0.9480	0.9743
cifar100	0.7683	0.7712	0.8435
dtd	0.6676	0.6984	0.7223
emnist_letters	0.9485	0.9334	0.9470
eurosat	0.9826	0.9711	0.9770
fashion_mnist	0.9353	0.9084	0.9282
fer2013	0.6553	0.5853	0.6495
food101	0.7303	0.7200	0.7763
gtsrb	0.9594	0.9174	0.9478
kmnist	0.9696	0.9188	0.9601
mnist	0.9946	0.9878	0.9946
oxford-iiit-pet	0.8613	0.9057	0.9201
oxford_flowers102	0.7637	0.7536	0.7291
pcam	0.8242	0.8357	0.8390
rendered-sst2	0.5211	0.5409	0.5041
resisc45	0.9063	0.8760	0.9030
stanford-cars	0.6241	0.5769	0.6260
stl10	0.9408	0.9668	0.9813
sun397	0.4794	0.5329	0.5742
svhn	0.9501	0.9262	0.9458

C.3. Detailed Experimental Results

This section provides the complete set of visualizations for the experimental results discussed in Section 5.2. Each figure corresponds to a specific hyperparameter analysis, illustrating the performance of the merged model across three different ResNet backbones (ResNet-18, ResNet-50, and ResNet-152). For each hyperparameter, we present both line plots showing trends with error bars (top row) and bar charts for direct comparison at specific values (bottom row), evaluated by both average accuracy and average loss.

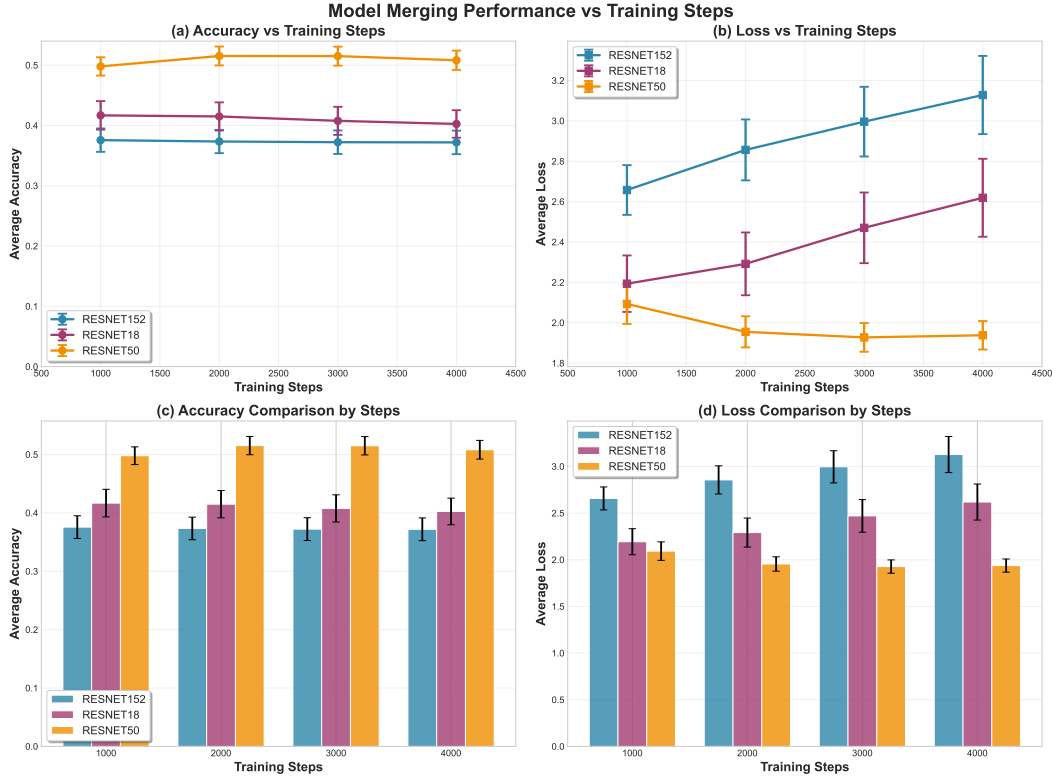


Figure 4. **Detailed results for the impact of fine-tuning steps (K_i).** The top row (a, b) shows performance trends, while the bottom row (c, d) provides direct comparisons. The results exhibit a clear trade-off: performance initially improves with more steps but then degrades, particularly for ResNet-18 and ResNet-50. This inverted U-shaped trend for accuracy and U-shaped trend for loss strongly substantiates our theoretical prediction of a balance between optimization and generalization, as discussed in Section 5.2.1.

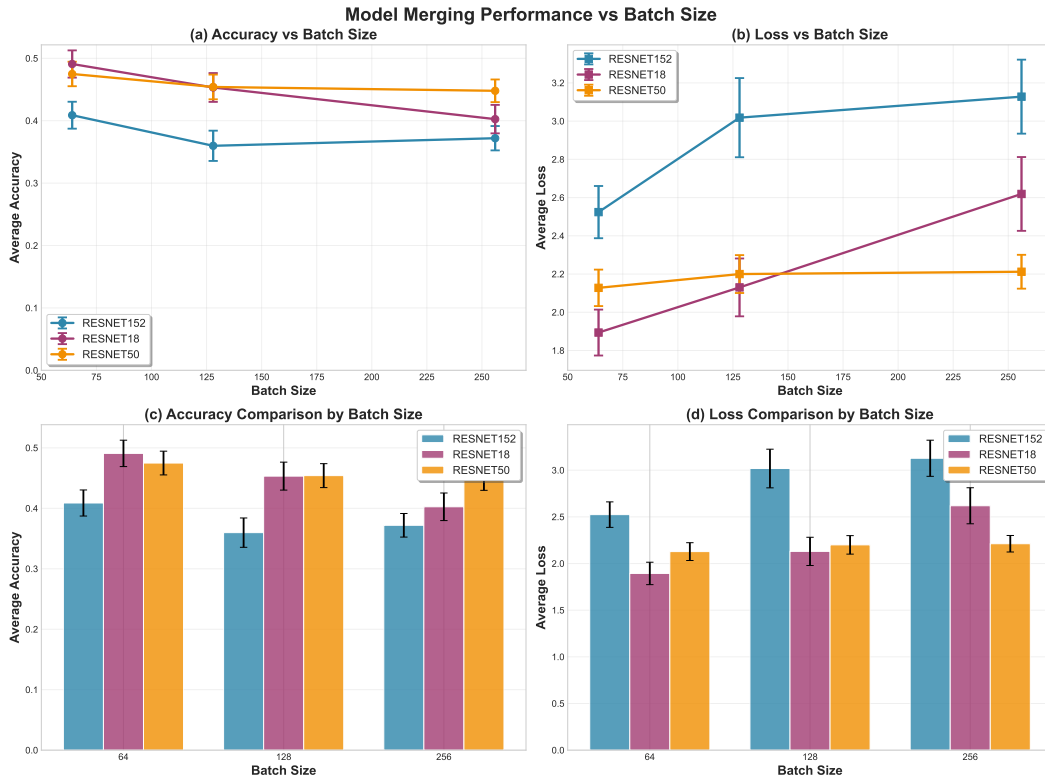


Figure 5. **Detailed results for the impact of batch size (b_i).** The plots show that for all three backbones, increasing the batch size from 64 to 256 leads to an improvement in performance, as evidenced by decreasing average accuracy (a) and increasing average loss (b). This suggests that in our experimental setting, the potential benefits of reduced gradient variance—a key component of our theoretical bound—outweigh the negative impact on model stability from larger batches, as discussed in Section 5.2.2.

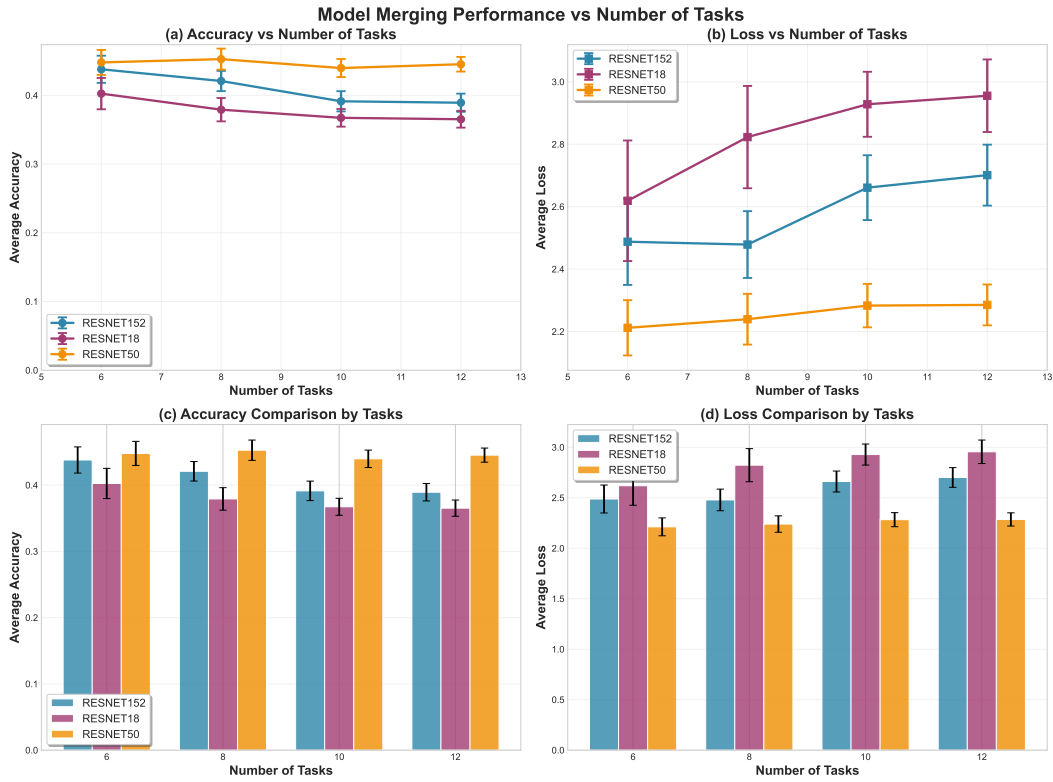


Figure 6. **Detailed results for the impact of the number of merged tasks (N).** As predicted by our theoretical analysis, increasing the number of tasks leads to a consistent decline in generalization performance. This is shown by the decreasing accuracy trends (a) and increasing loss trends (b) across all ResNet backbones. The results confirm that the penalty from accumulating task heterogeneity is the dominant factor, overwhelming the marginal benefits of optimization, as discussed in Section 5.2.5.

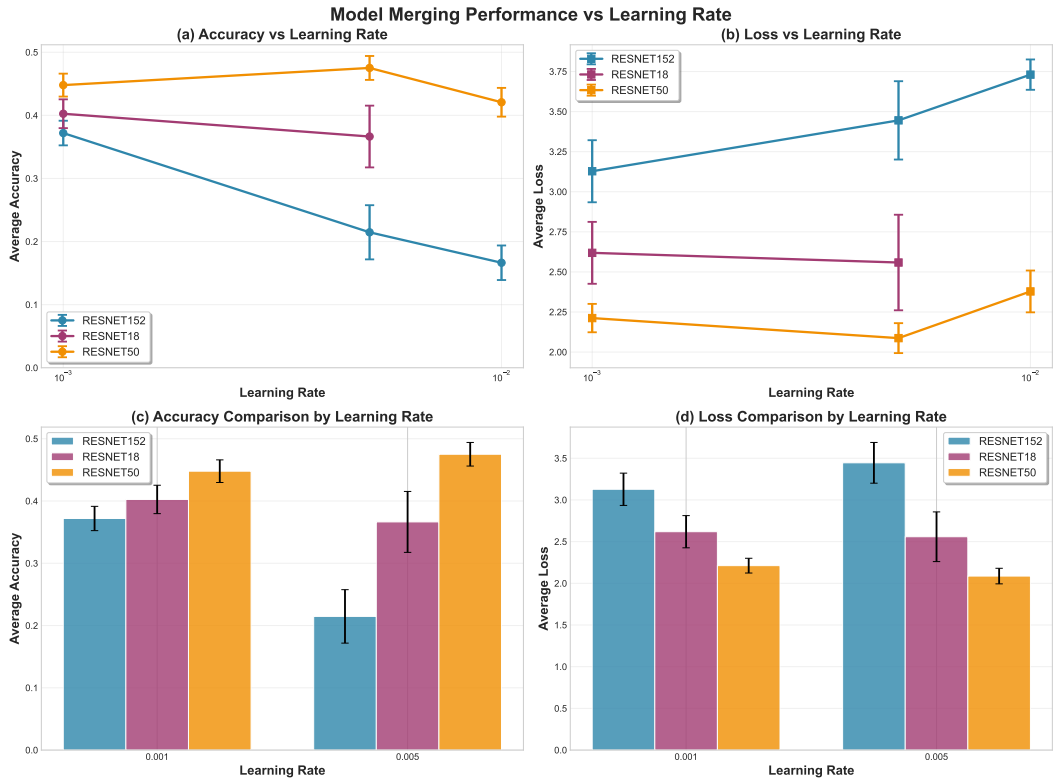


Figure 7. **Detailed results for the impact of the learning rate (η).** The plots demonstrate the critical role of the learning rate in model merging. A larger learning rate leads to a severe degradation in performance, with sharply decreasing accuracy (a) and increasing loss (b). This provides strong empirical evidence for our theoretical claim that the learning rate is a dominant factor controlling model stability, with larger values causing the stability component of our error bound to explode, as discussed in Section 5.2.3

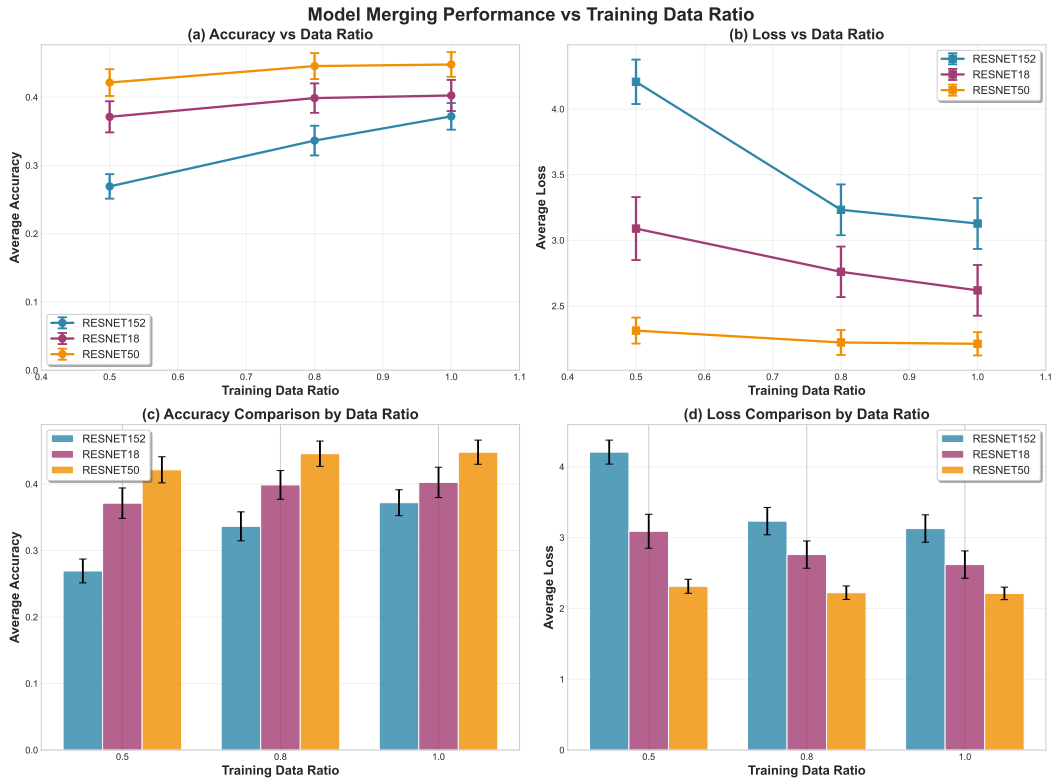


Figure 8. **Detailed results for the impact of the training data ratio (α_i).** The results show a clear and monotonic improvement in performance as more training data is used. Across all backbones, average accuracy (a) increases and average loss (b) decreases with a larger data ratio. This aligns perfectly with our theoretical prediction that more data enhances model stability and strictly tightens the generalization error bound, providing a "free lunch" for generalization, as discussed in Section 5.2.4.