

M-Loss: Quantifying Model Merging Compatibility with Limited Unlabeled Data

Tiantong Wang¹, Yiyang Duan¹, Haoyu Chen^{1,2}, Tiantong Wu^{3,*}, Wei Yang Bryan Lim^{1,*}

¹College of Computing and Data Science, Nanyang Technological University

²School of Computer and Information Technology, Beijing Jiaotong University

³Alibaba-NTU Global e-Sustainability CorpLab (ANGEL)

tiantong001@e.ntu.edu.sg, tiantong.wu@ntu.edu.sg, bryan.limwy@ntu.edu.sg

Abstract

Training of large-scale models is both computationally intensive and often constrained by the availability of labeled data. Model merging offers a compelling alternative by directly integrating the weights of multiple source models without requiring additional data or extensive training. However, conventional model merging techniques, such as parameter averaging, often suffer from the unintended combination of non-generalizable features, especially when source models exhibit significant weight disparities. Comparatively, model ensembling generally provides more stable and superior performance that aggregates multiple models by averaging outputs. However, it incurs higher inference costs and increased storage requirements. While previous studies experimentally showed the similarities between model merging and ensembling, theoretical evidence and evaluation metrics remain lacking. To address this gap, we introduce Merging-ensembling loss (*M-Loss*), a novel evaluation metric that quantifies the compatibility of merging source models using very limited unlabeled data. By measuring the discrepancy between parameter averaging and model ensembling at layer and node levels, *M-Loss* facilitates more effective merging strategies. Specifically, *M-Loss* serves both as a quantitative criterion of the **theoretical feasibility** of model merging, and a guide for **parameter significance** in model pruning. Our theoretical analysis and empirical evaluations demonstrate that incorporating *M-Loss* into the merging process significantly improves the alignment between merged models and model ensembling, providing a scalable and efficient framework for accurate model consolidation.

Code — <https://github.com/languangduan/mLoss>

1 Introduction

In large-scale deep learning, the reliance on large-scale labeled datasets and intensive computation limits the applicability of supervised methods (Cottier et al. 2024; Bang et al. 2024; Manchanda et al. 2024; Rakaraddi et al. 2024; Wang et al. 2021; Farboschi et al. 2024; Guo 2024). Model merging offers a promising alternative by directly fusing the weights of multiple pretrained or fine-tuned models into a single network, thereby reducing the cost for data collection and training (Jin et al. 2022; Li et al. 2023; Yadav

et al. 2024). It is particularly effective when labeled data is scarce or unavailable (Yang et al. 2023; Xu et al. 2024; Li et al. 2025). Additionally, merging multiple models often enhances performance, as the merged model benefits from the diverse perspectives and learned representations of its individual source models.

Despite its potential, model merging faces several challenges. Traditional methods averaging parameters without pre-processing can merge non-generalizable features (Cade 2015; Maleki et al. 2022). In scenarios where source models exhibit significant weight disparities, direct parameter averaging can yield poor results (Li et al. 2021). **Notably, existing research lacks a theoretical tool to assess model merging compatibility without labeled data.**

A theoretical challenge in model merging is understanding why and when multiple non-linear models can be effectively merged through simple weight averaging. While prior research has explored the feasibility of model merging by drawing parallels with model ensembling—such as Mode Connectivity (MC) and Feature Connectivity (FC) studies (Freeman and Bruna 2016; Garipov et al. 2018; Draxler et al. 2018; Zhou et al. 2023) – these studies primarily rely on empirical observations or theoretical analyses under specific conditions.

Model ensembling (Dash and Cooper 2004; Hoeting et al. 1999; Lv et al. 2023; Sagi and Rokach 2018; Wen, Tran, and Ba 2020) aggregates the outputs of multiple source models through linear averaging. It generally outperforms parameter averaging by providing more stable and accurate predictions with concrete theoretical evidence. However, ensembling necessitates storing multiple sets of model parameters and performing inference across all models, leading to increased computational and storage overhead.

To quantify the gap between model merging and model ensembling, we propose *M-Loss*, an evaluation metric designed to quantify the mergeability of source models. *M-Loss* measures the discrepancy between the outputs of a weight-merged model and those of an ensembled model using only unlabeled data. *M-Loss* can be applied to multiple scenarios to assess the mergeability of source models theoretically and predict the performance of the merged model without a labeled test set. *M-Loss* can also work as a reference for developing new merging methods. To make the problem setting concrete and to clarify how *M-Loss* in-

*Corresponding Authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

interfaces with practical merging algorithms, we provide an overview in Figure 1. It summarizes (i) how M-Loss quantifies the discrepancy between parameter averaging and ensembling using only unlabeled data, and (ii) how the resulting scores guide row-wise pruning schedules that plug into standard merging backends. In summary, our contributions are:

1. **M-Loss provides theoretical justification for model merging:** We address a fundamental question in model merging research: Why does a merged model, obtained through simple parameter averaging, closely approximate the ensembling of source models, despite non-linearity in model architecture? We propose a new evaluation metric called M-Loss, and derive the expectation of M-Loss under common activation functions, namely ReLU (Nair and Hinton 2010), GELU (Hendrycks and Gimpel 2016), and Leaky ReLU (Maas et al. 2013). Our analysis provides theoretical justification for why fine-tuned models originated from a shared pretrained backbone can be effectively merged.
2. **M-Loss as a reference for hyperparameter selection:** We show that M-Loss can act as a reliable indicator of parameter significance and inter-model conflicts, guiding the hyperparameter selection in model merging techniques such as TIES. By prioritizing parameters with lower M-Loss for retention, we propose the M-TIES method that employs a dynamical parameter pruning schedule for the merging process.

2 Preliminaries and Related Works

2.1 Model Merging

Model merging aims to consolidate multiple trained models into a single unified model by directly combining their parameters. Let $\{\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(n)}\}$ denote the parameters of n source models. A common approach of model merging is **simple averaging** (Dash and Cooper 2004; Hoeting et al. 1999; Wortsman et al. 2022), where the merged model θ_{merged} is computed as: $\theta_{\text{merged}} = \frac{1}{n} \sum_{i=1}^n \theta^{(i)}$, which implicitly assumes linear mode connectivity between models (Garipov et al. 2018). However, this simplistic approach faces great challenges under divergent optimization trajectories or when merging models fine-tuned for different tasks (Li et al. 2019). To address this, recent methods like **Task Arithmetic** (Ilharco et al. 2022) merge fine-tuned models by treating their parameter updates as additive task vectors relative to a shared pretrained initialization $\theta_{\text{pretrained}}$ as: $\theta_{\text{merged}} = \theta_{\text{pretrained}} + \sum_{i=1}^n \frac{\lambda}{n} (\theta^{(i)} - \theta_{\text{pretrained}})$, with tunable hyperparameter λ scaling the task vectors for better task adaptation. **AdaMerging** (Yang et al. 2023) with $\theta_{\text{merged}} = \sum_{i=1}^n \alpha_i \theta^{(i)}$ learns the optimal merging weights $\{\alpha_i\}_{i=1}^n$ by output entropy. However, multiple rounds of merging are required to obtain the optimal weights, which increases the computational cost.

2.2 Model Ensembling

Model ensembling (Dash and Cooper 2004; Hoeting et al. 1999; Lv et al. 2023; Sagi and Rokach 2018; Wen, Tran,

and Ba 2020) aggregates predictions from multiple models rather than their parameters. For input x , the ensemble output $f_{\text{ens}}(x)$ is typically a weighted average of individual model predictions. For equally ensembled n models:

$$f_{\text{ens}}(x) = \frac{1}{n} \sum_{i=1}^n f_{\theta^{(i)}}(x),$$

where $f_{\theta^{(i)}}(x)$ denotes the output of the i -th model. Ensembling leverages the ‘‘wisdom of the crowd’’ effect, reducing variance and improving robustness (Sagi and Rokach 2018). While ensembling outperforms naive parameter averaging (Lv et al. 2023), it requires storing all models and computing their joint outputs, which is infeasible for large-scale deployments.

2.3 Parameter Pruning in Model Merging

Task vectors from different fine-tuned models could have direction conflicts, degrading the merged model’s performance. Pruning resolves parameter conflicts during merging by selectively discarding weights in task vectors. Let the task vector from the i -th model be $\Delta^{(i)} = \theta^{(i)} - \theta_{\text{pretrained}}$. **TIES** (Yadav et al. 2024) consists of three key stages: trim redundant parameters, elect dominant signs, and merge the aligned parameters. **DARE** (Yu et al. 2024) shows the importance of rescaling after discarding. It replaces step 1 in TIES with randomly keeping $k\%$ parameters and rescaling the kept parameters by $1/k\%$. However, DARE requires the model to be large so that random trimming would not incur feature loss.

Our work introduces M-Loss to guide pruning by directly measuring functional discrepancies between merged and ensemble predictions, eliminating reliance on additional labeled data.

3 Problem Formulation

3.1 Problem Description

Consider a set of models $\{M_1, M_2, \dots, M_q\}$ that share the same architecture, consisting of $L + 1$ layers with input dimension m_0 , hidden dimensions m_1, \dots, m_L , and output dimension m_{L+1} . The neural network function for model M is defined as:

$$f_M(x) = W_M^{L+1} \sigma(H_M^L \dots \sigma(H_M^1(x))) + b_M^{L+1}, \quad (1)$$

where $H_M^k(x) = W_M^k x + b_M^k$, for $k = 1, \dots, L$, represents the linear transformation of the k -th layer. Here, W_M^k and b_M^k denote the weights and biases of the k -th layer of model M , respectively.

The pre-activation vector of the j -th layer is defined as:

$$h_M^j(x) = H_M^j \sigma(H_M^{j-1} \dots \sigma(H_M^1(x))), \quad (2)$$

for $j = 1, \dots, L$. We write h_j for simplicity if M and x are clear from context. The i -th entry is written as $h_M^{j,i}(x)$ or $h_{j,i}$.

The discrepancy between model merging and model ensembling arises mainly from the non-linear activation functions. To address this, we investigate the flow of intermediate representations around the activation functions.

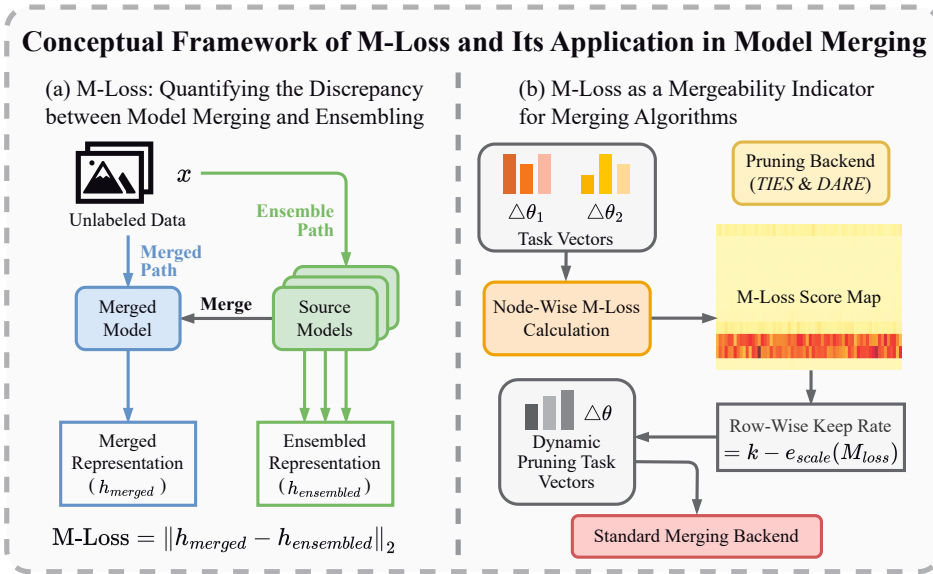


Figure 1: Conceptual overview of M-Loss and its use in M-TIES. (a) M-Loss measures the discrepancy between parameter-averaged and ensembled representations on unlabeled data, producing layer-/node-wise scores. (b) The node-wise M-Loss score map drives dynamic row-wise keep rates, which integrate with standard merging backends (e.g., TIES Top-K or DARE) to improve mergeability and efficiency.

3.2 Definitions

To establish our M-Loss theory, we first need to introduce the following definitions:

Definition 3.1 (Intermediate Representation) The intermediate representation of a given input vector x with respect to a model M refers to the representation formed by partially passing the input vector x through the model (e.g., pre-/post-activation vectors h_k, x_k , their entries $h_{k,i}, x_{k,i}$, the input and output of hidden layers as shown in Figure 2).

For example, the intermediate representation h_k passes through the activation function in layer k to obtain the post-activation intermediate representation x_k .

Definition 3.2 (Linearly Correlated Model Parameters)

For any intermediate representation $h_{k,i}$ of a model M characterized by parameter set Θ , the Linearly Correlated Model Parameters (LCP) of $h_{k,i}$ (or of node $N_{k,i}$) are defined as the parameter subset holding constant partial derivative determined by input vector only: $\{w \in \Theta \mid \frac{\partial h_{k,i}}{\partial w} = C, \text{ where } C \text{ is a nonzero constant given by input vector } x\}$.

For example, given a simple layer without bias, denoted as $h_k = Wx$, the LCP with respect to $h_{k,i}$ is the parameters in the i -th row of W because $\frac{\partial h_{k,i}}{\partial W_{i,j}} = x_j$. An example is shown in Figure 2.

Definition 3.3 (M-Loss on Layer Level) For a given input vector x , pretrained models $\mathcal{M} = \{M_1, \dots, M_q\}$, merging weights $\alpha = (\alpha_1, \dots, \alpha_q)$, and the j -th layer of model M , the merging-ensembling loss (M-Loss) is defined as:

$$h = (h_{M_1}^j(x), \dots, h_{M_q}^j(x)), \quad \alpha \cdot h = \sum_{p=1}^q \alpha_p h_{M_p}^j(x).$$

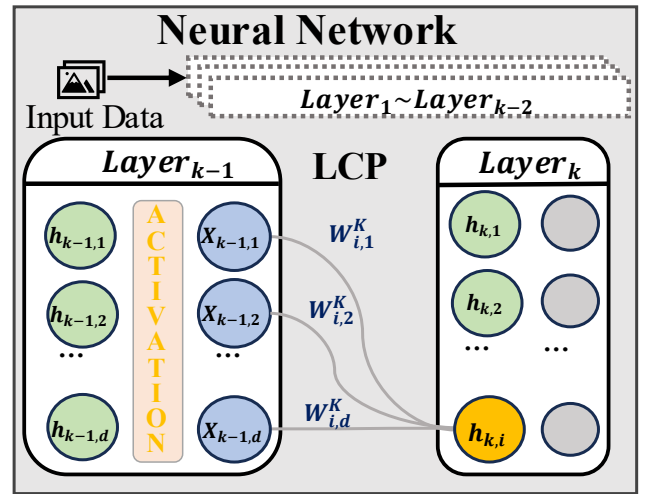


Figure 2: Visualization of Linearly Correlated Parameters (LCP) in a neural network. The figure illustrates how the partial derivative $\frac{\partial h_{k,i}}{\partial W_{i,j}^k}$ relates to the input from the previous layer $x_{(k-1),j}$.

$$\mathcal{L}(W_{\mathcal{M}}^j)(x) = \|\sigma(\alpha \cdot h) - \alpha \cdot \sigma(h)\|_2 \quad (3)$$

where $\|\cdot\|_2$ denotes the L_2 -norm.

Remark 1 If the merging weights α_i are not specified, they are set to $\frac{1}{q}$ by default for simple averaging.

Definition 3.4 (M-Loss on Node Level) For a given input vector x , pretrained models $\mathcal{M} = \{M_1, \dots, M_q\}$, merging weights $\alpha = (\alpha_1, \dots, \alpha_q)$, and the i -th node of the k -

th layer in model M , the M-Loss is defined as the absolute value of output discrepancy:

$$h = (h_{M_1}^{k,i}(x), \dots, h_{M_q}^{k,i}(x)), \quad \alpha \cdot h = \sum_{p=1}^q \alpha_p h_{M_p}^{k,i}(x).$$

$$\mathcal{L}(N_{k,i}^{\mathcal{M}})(x) = \left| \sigma(\alpha \cdot h) - \alpha \cdot \sigma(h) \right| \quad (4)$$

We write $\mathcal{L}(N_{k,i})$ for simplicity if \mathcal{M} is clear from context.

Remark 2 (Relation between Node and Layer M-Loss)

Note that by definition, we have:

$$\mathcal{L}(W_{\mathcal{M}}^k)(x) = \sqrt{\sum_{i=1}^d \mathcal{L}(N_{k,i}^{\mathcal{M}})^2}, \quad (5)$$

where d is the hidden dimension of the model.

For models without layer normalization (Ba, Kiros, and Hinton 2016), we adopt the following normalized M-Loss to avoid the effect of scaling among layers.

Definition 3.5 (Normalized M-Loss on Layer Level) For a given input x , pretrained models $\mathcal{M} = \{M_1, \dots, M_q\}$ and merging weights $\alpha = (\alpha_1, \dots, \alpha_q)$, denote

$$h = (h_{M_1}^j(x), \dots, h_{M_q}^j(x)), \quad \alpha \cdot h = \sum_{p=1}^q \alpha_p h_{M_p}^j(x).$$

Then the normalized merging-ensembling loss is

$$\mathcal{L}_{\text{norm}}(W_{\mathcal{M}}^j)(x) = \frac{\|\sigma(\alpha \cdot h) - \alpha \cdot \sigma(h)\|_2}{\|\sigma(\alpha \cdot h)\|_2 + \epsilon}, \quad (6)$$

where $\epsilon > 0$ is a small constant added in the denominator to ensure numerical stability (typically $\epsilon = 10^{-4}$ or smaller).

Definition 3.6 (Normalized M-Loss on Node Level)

For a given input vector x , pretrained models $\mathcal{M} = \{M_1, \dots, M_q\}$ and merging weights $\alpha = (\alpha_1, \dots, \alpha_q)$, denote

$$h = (h_{M_1}^{k,i}(x), \dots, h_{M_q}^{k,i}(x)), \quad \alpha \cdot h = \sum_{p=1}^q \alpha_p h_{M_p}^{k,i}(x).$$

Then the normalized node-level merging-ensembling loss is

$$\mathcal{L}_{\text{norm}}(N_{k,i}^{\mathcal{M}})(x) = \frac{|\sigma(\alpha \cdot h) - \alpha \cdot \sigma(h)|}{|\sigma(\alpha \cdot h)| + \epsilon}, \quad (7)$$

where $\epsilon > 0$ is a small constant added in the denominator to ensure numerical stability (typically $\epsilon = 10^{-4}$ or smaller).

Model merging operates on weights rather than alignment through non-linear layers. Empirical results (Jin et al. 2022; Dai et al. 2025) indicate that linear layers are more significant to the merging process than non-linear layers.

4 M-Loss as a Quantitative Criterion for Theoretical Merging Feasibility

The central question on model merging feasibility is: why can simple parameter averaging approximate output-level ensembling for non-linear networks? In this section, we address this challenge under the M-Loss theory.

By Remark 2, layer-level M-Loss is from the accumulation of node-level M-Loss. Meanwhile, the discrepancy between model merging and ensembled output is from the accumulation of layer-level M-Loss, as the discrepancy is accumulated when passing through each activation function. Elementally, we calculate the expected node-level M-Loss to bound the discrepancy between the model merging output and the ensembled output.

Consider the activation function $f(\cdot)$. We aim to compute the expected M-Loss on the node level as:

$$\mathbb{E}[D] = \mathbb{E}_{x,a,b} \left[\left| f\left(\frac{a+b}{2}\right) - \frac{f(a)+f(b)}{2} \right| \right],$$

where:

- $x \sim \text{Uniform}(-k, k)$: Approximates the **intermediate representation** of the pre-activation value through the pretrained model.
- $a, b \sim \mathcal{N}(x, \sigma^2)$: Approximate **intermediate representation** of pre-activation values through fine-tuned models, which correspond to two different $h_{k,i}$ in the previous section
- Assume that $k \gg \sigma$, as fine-tuned models have a small amount of parameter shift from the pre-trained model.

For the three commonly used activation functions: ReLU (Nair and Hinton 2010), GELU (Hendrycks and Gimpel 2016), and Leaky ReLU (Maas et al. 2013), we have the following estimation (see Appendix for detailed derivations):

$$\mathbb{E}[D_{\text{ReLU}}] \approx \frac{\sigma^2}{\sqrt{2\pi k}}$$

$$\mathbb{E}[D_{\text{LeakyReLU}, \alpha}] \approx \frac{(1-\alpha)\sigma^2}{\sqrt{2\pi k}}$$

$$\mathbb{E}[D_{\text{GELU}}] \approx \frac{\sigma^2}{4k}$$

As k is the range of intermediate representation of the pretrained model, estimation of M-Loss indicates that merging models fine-tuned from a more general pretrained model (larger range k) would perform better. Meanwhile, σ measures the expected parameter difference of source models, indicating that mergeability is positively correlated with parameter similarity.

When the assumption $k \gg \sigma$ holds, the expectations calculated above are all small, so the expectation of M-Loss for merging fine-tuned models from the same pretrained model is small. Thus, we theoretically proved that the merged models obtained by averaging parameters behave similarly to averaging the outputs of each individual model.

Algorithm 1: M-TIES Merging

```
1: Input: Pretrained model  $M$  with parameters  $\theta$ 
   consisting of  $L$  layers  $\theta^{(1)}, \dots, \theta^{(L)}$ ; source models
    $M_1, \dots, M_q$  with parameters  $\theta_1, \dots, \theta_q$ ; merging
   weights  $\alpha_1, \dots, \alpha_q$ ; unlabeled dataset  $\mathcal{D}$ ; base keep per-
   centage  $k\%$ ; pruning variation  $e\%$ .
2: Output: Merged model  $M_{\text{merged}}$ .
3: Initialize  $M_{\text{merged}} \leftarrow M$ .
4: for  $l = 1$  to  $L$  do
5:   # Compute task vectors for layer  $l$ 
6:   for  $i = 1$  to  $q$  do
7:      $T_i^{(l)} \leftarrow \theta_i^{(l)} - \theta^{(l)}$ 
8:   end for
9:   # Compute node-wise M-Loss for layer  $l$ 
10:  for  $j = 1$  to  $d_l$  do
11:    Compute  $\mathcal{L}(N_{l,j})$  on layer  $l$ .
12:  end for
13:  # Determine keep ratios based on M-Loss
14:  for  $j = 1$  to  $d_l$  do
15:     $k_{l,j} \leftarrow \text{RankNorm}(\mathcal{L}(N_{l,j})) \in [(k - e)\%, k\%]$ 
16:  end for
17:  # Prune task vectors at node level
18:  for  $j = 1$  to  $d_l$  do
19:    Prune parameters in  $\{T_i^{(l)}\}_{i=1}^q$  linearly correlated
    with  $N_{l,j}$ , keeping proportion  $k_{l,j}$ .
20:  end for
21:  # Elect sign and merge pruned task vectors
22:   $\Delta\theta^{(l)} \leftarrow \sum_{i=1}^q \alpha_i \text{ElectPruned}(T_i^{(l)})$ 
23:  Update layer  $l$ :  $\theta^{(l)} \leftarrow \theta^{(l)} + \Delta\theta^{(l)}$  in  $M_{\text{merged}}$ .
24:  # Synchronize merged layer to source models
25:  for  $i = 1$  to  $q$  do
26:     $\theta_i^{(l)} \leftarrow \theta^{(l)}$ 
27:  end for
28: end for
29: return  $M_{\text{merged}}$ .
```

5 M-Loss as an Indicator of Hyperparameter Selection

We show that M-Loss can act as an indicator of parameter conflicts between models, guiding the pruning budget schedule of other merging methods such as TIES. By prioritizing parameters with lower M-Loss for retention, we develop a dynamic pruning rate scheduler as a plug-in to the existing merging method TIES (Yadav et al. 2024).

Different from the TIES method with a fixed trimming rate, we propose an enhanced pruning strategy inspired by TIES merging, leveraging the M-Loss metric. We determine the pruning (trimming) proportion in a more dedicated way, computed from the M-Loss.

M-TIES Merging receives a pretrained base model M (parameters θ), q source models M_1, \dots, M_q (parameters $\theta_1, \dots, \theta_q$), merging weights $\alpha_1, \dots, \alpha_q$, an unlabeled dataset \mathcal{D} , a base keep-rate $k\%$ and a pruning variation $e\%$. It then processes each layer $l = 1, \dots, L$ in turn: first computing task vectors $T_i^{(l)} = \theta_i^{(l)} - \theta^{(l)}$ for all sources;

next evaluating a node-level M-Loss $\mathcal{L}(N_{l,j})$ on \mathcal{D} ; rank-normalizing these losses into $[(k - e)\%, k\%]$ to obtain per-node keep ratios $k_{l,j}$; pruning each $T_i^{(l)}$ by zeroing out the smallest in absolute value $(100 - k_{l,j})\%$ of parameters linearly correlated with node j ; then elect the dominant sign of each entry, forming a weighted merge increment $\Delta\theta^{(l)} = \sum_i \alpha_i \text{ElectPruned}(T_i^{(l)})$ which is added back to the base layer: $\theta^{(l)} \leftarrow \theta^{(l)} + \Delta\theta^{(l)}$. Finally, the updated $\theta^{(l)}$ is synchronized to all sources before proceeding to the next layer. After all layers are merged, the resulting model M_{merged} is returned. This method effectively balances the integration of multiple models while selectively pruning redundant or conflicting parameters, ensuring an optimal trade-off between model performance and parameter efficiency.

6 Experiments

6.1 Experimental Setup

To validate the effectiveness of the proposed M-TIES method, we conduct model merging on Vision Transformer (ViT) (Alexey 2020) with pretrained OpenAI CLIP (Radford et al. 2021) ViT-B/32 and ViT-L-14 models. We use the open-source model checkpoints fine-tuned on eight datasets of diverse types of tasks – RESISC45 (Cheng, Han, and Lu 2017), Cars (Krause et al. 2013), MNIST (LeCun et al. 1998), DTD (Cimpoi et al. 2014), EuroSAT (Helber et al. 2019), GTSRB (Stallkamp et al. 2012), SUN397 (Xiao et al. 2010) and SVHN (Netzer et al. 2011) as source models. The aggregation weights α_i 's are set to be equal. We compare our M-TIES method with other model merging baselines, namely **Simple Average** (Dash and Cooper 2004; Hoeting et al. 1999; Wortsman et al. 2022), **Task Arithmetic** (Ilharco et al. 2022), **TIES** (Yadav et al. 2024) and **DARE** (Yu et al. 2024). We discuss in the Appendix for baseline selection. For the unlabeled dataset selection, we only use 128 samples in total, which means 16 samples for each dataset on average. This is a very few-shot setting and is practical in real-world scenarios. For other layers, as they are less important in merging (Jin et al. 2022; Dai et al. 2025) and require a more complex procedure for computing M-Loss than linear layers, we adopt vanilla trimming with TIES. We conduct all experiments using a single NVIDIA RTX A6000 GPU.

6.2 Experimental Results

The experimental results for our method, four merging baselines, and ensembling are shown in Table 1, covering experiments on both the ViT-B/32 and the larger ViT-L/14 model. For each method except for simple averaging, we first conduct a hyperparameter search on the validation set using the ViT-B/32 model and report the best average accuracy. The hyperparameter search range is included in the Appendix. We then repeat the experiments on the ViT-L/14 model using these same fixed hyperparameters. As explained in the table's caption, for each backbone, numbers in **bold** are the top accuracies for the task, and underlined numbers denote the second best.

| Backbone | Method | RESISC45 | Cars | MNIST | DTD | EuroSAT | GTSRB | SUN397 | SVHN | Avg |
|----------|-----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| ViT-B/32 | M-TIES | <u>72.60</u> | 61.07 | 97.62 | <u>54.84</u> | <u>82.02</u> | <u>72.44</u> | 62.19 | 83.06 | <u>73.23</u> |
| | TIES | 70.67 | 58.61 | 98.30 | 54.20 | 80.22 | 72.11 | 59.01 | 86.20 | 72.42 |
| | Task Arithmetic | 71.27 | 60.70 | 95.32 | 51.76 | 79.74 | 67.32 | 62.06 | 76.68 | 70.61 |
| | Simple Avg. | 71.46 | <u>63.34</u> | 87.46 | 50.11 | 73.00 | 52.79 | <u>64.91</u> | 64.16 | 65.90 |
| | DARE | 69.97 | <u>57.98</u> | <u>97.95</u> | 53.24 | 78.89 | 72.00 | <u>59.14</u> | <u>83.96</u> | 71.64 |
| | Ensembling | 79.87 | 66.60 | 95.80 | 58.30 | 98.30 | 81.11 | 66.35 | <u>82.15</u> | 78.56 |
| ViT-L/14 | M-TIES | 88.57 | <u>83.35</u> | 99.06 | 66.91 | <u>94.61</u> | 83.80 | 76.13 | 89.78 | 85.28 |
| | TIES | 88.19 | 82.81 | <u>99.01</u> | 66.70 | 94.37 | 83.36 | <u>75.65</u> | <u>89.42</u> | 84.94 |
| | Task Arithmetic | 86.17 | 82.44 | 98.54 | 65.59 | 93.93 | 83.47 | <u>73.56</u> | <u>85.26</u> | 83.62 |
| | Simple Avg. | 82.67 | 81.54 | 97.01 | 62.77 | 91.17 | 70.63 | 71.65 | 78.23 | 79.46 |
| | DARE | <u>88.33</u> | <u>83.35</u> | 98.97 | <u>66.86</u> | 94.06 | <u>84.20</u> | 75.37 | 89.19 | 85.04 |
| | Ensembling | 87.73 | 85.36 | 98.78 | 66.81 | 98.24 | 87.92 | 74.76 | 84.92 | 85.56 |

Table 1: Accuracy comparison of merging methods on ViT-B/32 and ViT-L/14 backbones. For each backbone and each dataset, the best result is in bold and the second-best is underlined. M-TIES consistently outperforms other merging methods and achieves competitive results against the much more costly Ensembling baseline, especially on the larger ViT-L/14 model.

| Sample Size | RESISC45 | Cars | MNIST | DTD | EuroSAT | GTSRB | SUN397 | SVHN | Avg |
|-------------|----------|--------|--------|--------|---------|--------|--------|--------|--------|
| 128 | 72.603 | 61.074 | 97.620 | 54.840 | 82.019 | 72.439 | 62.195 | 83.063 | 73.232 |
| 256 | 72.571 | 61.074 | 97.620 | 54.840 | 82.000 | 72.462 | 62.186 | 83.059 | 73.227 |
| 512 | 72.571 | 61.074 | 97.620 | 54.840 | 82.019 | 72.454 | 62.176 | 83.059 | 73.227 |

Table 2: Accuracy of M-TIES merging ViT-B/32 fine-tuned models with different sample sizes for computing M-loss.

| Random Seed | RESISC45 | Cars | MNIST | DTD | EuroSAT | GTSRB | SUN397 | SVHN | Avg |
|-------------|----------|--------|-------|--------|---------|--------|--------|--------|--------|
| 1 | 72.603 | 61.074 | 97.62 | 54.840 | 82.278 | 72.478 | 61.855 | 83.063 | 73.225 |
| 2 | 72.587 | 61.099 | 97.61 | 54.840 | 81.611 | 72.486 | 62.025 | 83.059 | 73.165 |
| 42 | 72.603 | 61.074 | 97.62 | 54.840 | 82.019 | 72.439 | 62.195 | 83.063 | 73.232 |
| Std | 0.009 | 0.014 | 0.006 | 0.000 | 0.336 | 0.025 | 0.170 | 0.002 | 0.037 |

Table 3: Accuracy and standard deviation of M-TIES merging ViT-B/32 fine-tuned models under different random seeds.

6.3 Experiment Results Analysis

General Comparison From Table 1, we can see that our proposed M-TIES method outperforms other merging baselines in both ViT-B/32 and ViT-L/14 model backbones, while below that of model ensembling. This verifies our motivation to bridge the gap between merging and ensembling by M-Loss. Meanwhile, calculation gives the variance of accuracy for top-3 merging methods M-TIES, TIES, and DARE are 172.22 and 203.27. and 197.61 (in ViT-B/32) respectively, which shows that M-TIES holds better stability and fairness (not favoring high-accuracy tasks). Additionally, when the accuracy gap between merging and ensembling is small, M-TIES even outperforms ensembling for certain tasks.

Computational Cost Analysis As in our algorithm, when merging layer k , the previous $(k - 1)$ models would have been updated to the same merged weight, so we only need to pass the input to the previous layer once and get the shared input of layer k , then compute the M-Loss for nodes in layer k . The inference cost is passing input to $(k - 1 + q)$ layers, which is much less than passing the input to each model in-

dividually, which needs kq layers (q is the number of source models).

For the empirical time consumption, M-TIES does not add a significant time cost. M-TIES uses approximately 1 minute and 30 seconds and 3 minutes for merging ViT-B/32 and ViT-L/14 models, respectively, compared with the TIES method, which takes 30 seconds and 1 minute, as the computation of M-Loss only involves forward inference of a small unlabeled dataset. The main time cost for the experiment is the evaluation process, which takes around 5 minutes and 15 minutes for ViT-B/32 and ViT-L/14, respectively.

6.4 Sampling Efficiency and Stability over Randomness in Evaluating M-Loss

Table 2 shows that computing M-Loss only requires very limited unlabeled data. As we vary the sample size for evaluating the M-Loss, we find that there is little change in the final results. Thus, we chose only one batch of 128 samples in our experiments. For the randomness test, we run the experiments over three random seeds for drawing 128 samples as in Table 3, with a low standard deviation reported.

| Method | Flowers102 | FashionMNIST | Food101 | STL10 | CIFAR100 | Avg |
|------------|--------------|--------------|--------------|--------------|--------------|--------------|
| M-TIES | <u>54.82</u> | <u>62.94</u> | <u>61.02</u> | <u>89.90</u> | <u>48.60</u> | <u>63.46</u> |
| TIES | 48.40 | 58.61 | 52.35 | 87.60 | 40.78 | 57.55 |
| DARE | 51.86 | 62.28 | 55.14 | 87.80 | 44.19 | 60.25 |
| Ensembling | 62.77 | 67.93 | 72.61 | 96.00 | 61.70 | 72.20 |

Table 4: Accuracy comparison of different merging methods on out-of-domain datasets, using ViT-B/32 models.

| Layer Index | RESISC45 | Cars | MNIST | DTD | EuroSAT | GTSRB | SUN397 | SVHN | Avg |
|-------------|----------|--------|--------|--------|---------|--------|--------|--------|--------|
| All | 72.603 | 61.074 | 97.620 | 54.840 | 82.019 | 72.439 | 62.195 | 83.063 | 73.232 |
| 0,8,9,10 | 72.619 | 61.112 | 97.570 | 55.000 | 82.000 | 72.328 | 62.337 | 82.760 | 73.216 |
| 8,9,10 | 72.619 | 61.099 | 97.560 | 55.000 | 82.074 | 72.312 | 62.333 | 82.771 | 73.221 |

Table 5: Accuracy of M-TIES merging ViT-B/32 fine-tuned models with different layers pruned by M-loss.

| Layer Index | RESISC45 | Cars | MNIST | DTD | EuroSAT | GTSRB | SUN397 | SVHN | Avg |
|---------------|----------|--------|--------|--------|---------|--------|--------|--------|--------|
| All | 88.571 | 83.348 | 99.060 | 66.915 | 94.611 | 83.800 | 76.134 | 89.782 | 85.278 |
| 0, 20, 21, 22 | 88.587 | 83.273 | 99.070 | 66.809 | 94.593 | 83.903 | 76.088 | 89.747 | 85.259 |
| 20, 21, 22 | 88.571 | 83.273 | 99.070 | 66.809 | 94.611 | 83.895 | 76.065 | 89.751 | 85.256 |

Table 6: Accuracy of M-TIES merging ViT-L/14 fine-tuned models with different layers pruned by M-loss.

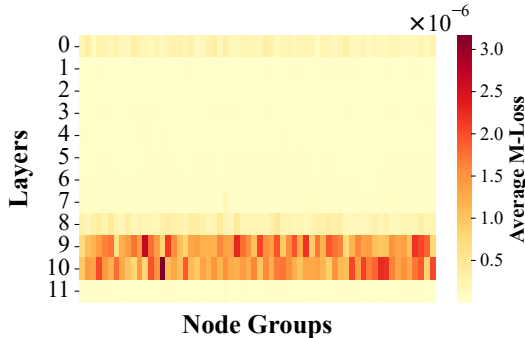


Figure 3: Layerwise node group M-Loss distribution across different layers of ViT-B/32 models. Each colored block reveals the average M-Loss of 50 consecutive nodes, with the x-axis being the node group number and the y-axis being the layer number.

It shows that our method holds great stability over random sampling.

6.5 Out-of-Domain Generalizability of M-TIES

We also conduct out-of-domain testing, by evaluating the ensemble model and top-3 ViT-B/32 models merged as in Table 1 on Flowers 102 (Nilsback and Zisserman 2008), FashionMNIST (Xiao, Rasul, and Vollgraf 2017), Food 101 (Bossard, Guillaumin, and Van Gool 2014), STL 10 (Coates, Ng, and Lee 2011), and CIFAR100 (Krizhevsky and Hinton 2009). As shown in Table 4, M-TIES outperforms other baselines except for ensembling, which shows its great generalizability on unseen tasks.

6.6 Few-Layer M-TIES: a More Efficient Alternative

We also analyze the M-Loss value across different nodes in different layers during the merging process in Figure 3. We find that the first layer and the last few layers (excluding the output layer) hold higher M-Loss than other layers.

With this insight, we propose Few-Layer M-TIES, which only uses a dynamic pruning method for the first and last few layers, so as to save the computational cost. Experiment results in Table 5 and Table 6 show that there is almost no change to the performance, but the computational cost is largely reduced.

7 Conclusion

This paper introduces **M-Loss**, a novel metric that quantifies the gap between parameter-averaged and output-averaged models without relying on labeled data. By computing the expected M-Loss for common activation functions, we provide theoretical justification for the conditions where parameter averaging yields predictions close to model ensembling, thereby establishing theoretical model merging feasibility. To show that M-Loss can be integrated with a concrete merging method, we integrate the M-Loss dynamic budget scheduler into the TIES merging framework, guiding the selective removal of conflicting parameters. The integration leads to superior performance compared to existing methods. Empirical evaluation results on ViT models underscore M-Loss’s effectiveness in identifying crucial parameters. Overall, this work advances the theoretical foundations of model merging and contributes practical tools for the efficient merging of multiple models.

Acknowledgments

This research is supported by the NTU startup grant and the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL). This research/project is supported by A*STAR under its Japan-Singapore Joint Call: JST-A*STAR 2024 (Project ID: R24I6IR139).

References

- Alexey, D. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *arXiv:1607.06450*.
- Bang, J.; Choi, Y.; Kim, M.; Kim, Y.; and Rhu, M. 2024. vtrain: A simulation framework for evaluating cost-effective and compute-optimal large language model training. In *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 153–167. IEEE.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101 – Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision*, volume 8694 of *Lecture Notes in Computer Science*, 446–461.
- Cade, B. S. 2015. Model averaging and muddled multimodel inferences. *Ecology*, 96(9): 2370–2382.
- Cheng, G.; Han, J.; and Lu, X. 2017. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10): 1865–1883.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; and Vedaldi, A. 2014. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3606–3613.
- Coates, A.; Ng, A.; and Lee, H. 2011. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, 215–223. Fort Lauderdale, FL, USA: PMLR.
- Cottier, B.; Rahman, R.; Fattorini, L.; Maslej, N.; and Owen, D. 2024. The rising costs of training frontier AI models. *arXiv preprint arXiv:2405.21015*.
- Dai, R.; Hu, S.; Shen, X.; Zhang, Y.; Tian, X.; and Ye, J. 2025. Leveraging Submodule Linearity Enhances Task Arithmetic Performance in LLMs. *arXiv preprint arXiv:2504.10902*.
- Dash, D.; and Cooper, G. F. 2004. Model averaging for prediction with discrete Bayesian networks. *Journal of Machine Learning Research*, 5(Sep): 1177–1203.
- Draxler, F.; Veschgini, K.; Salmhofer, M.; and Hamprecht, F. 2018. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*, 1309–1318. PMLR.
- Faraboschi, P.; Giles, E.; Hotard, J.; Owczarek, K.; and Wheeler, A. 2024. Reducing the Barriers to Entry for Foundation Model Training. *arXiv preprint arXiv:2404.08811*.
- Freeman, C. D.; and Bruna, J. 2016. Topology and geometry of half-rectified network optimization. *arXiv preprint arXiv:1611.01540*.
- Garipov, T.; Izmailov, P.; Podoprikin, D.; Vetrov, D. P.; and Wilson, A. G. 2018. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31.
- Guo, Z. 2024. More Compute Is What You Need. *arXiv preprint arXiv:2404.19484*.
- Helber, P.; Bischke, B.; Dengel, A.; and Borth, D. 2019. 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.
- Hendrycks, D.; and Gimpel, K. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Hoeting, J. A.; Madigan, D.; Raftery, A. E.; and Volinsky, C. T. 1999. Bayesian model averaging: a tutorial (with comments by M. Clyde, David Draper and El George, and a rejoinder by the authors. *Statistical science*, 14(4): 382–417.
- Ilharco, G.; Ribeiro, M. T.; Wortsman, M.; Gururangan, S.; Schmidt, L.; Hajishirzi, H.; and Farhadi, A. 2022. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*.
- Jin, X.; Ren, X.; Preotiuc-Pietro, D.; and Cheng, P. 2022. Dataless knowledge fusion by merging weights of language models. *arXiv preprint arXiv:2212.09849*.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, 554–561.
- Krizhevsky, A.; and Hinton, G. 2009. Learning Multiple Layers of Features from Tiny Images. Technical Report 0, University of Toronto, Toronto, Ontario.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Li, W.; Gao, H.-a.; Gao, M.; Tian, B.; Zhi, R.; and Zhao, H. 2025. Training-free model merging for multi-target domain adaptation. In *European Conference on Computer Vision*, 419–438. Springer.
- Li, W.; Peng, Y.; Zhang, M.; Ding, L.; Hu, H.; and Shen, L. 2023. Deep model fusion: A survey. *arXiv preprint arXiv:2309.15698*.
- Li, X.; Huang, K.; Yang, W.; Wang, S.; and Zhang, Z. 2019. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*.
- Li, X.; Jiang, M.; Zhang, X.; Kamp, M.; and Dou, Q. 2021. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*.
- Lv, X.; Ding, N.; Qin, Y.; Liu, Z.; and Sun, M. 2023. Parameter-efficient weight ensembling facilitates task-level

- knowledge transfer. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 270–282.
- Maas, A. L.; Hannun, A. Y.; Ng, A. Y.; et al. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, 3. Atlanta, GA.
- Maleki, F.; Ovens, K.; Gupta, R.; Reinhold, C.; Spatz, A.; and Forghani, R. 2022. Generalizability of machine learning models: quantitative evaluation of three methodological pitfalls. *Radiology: Artificial Intelligence*, 5(1): e220028.
- Manchanda, S.; Gupta, S.; Ranu, S.; and Bedathur, S. J. 2024. Generative modeling of labeled graphs under data scarcity. In *Learning on Graphs Conference*, 32–1. PMLR.
- Nair, V.; and Hinton, G. E. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, 807–814.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, 4. Granada.
- Nilsback, M.-E.; and Zisserman, A. 2008. Automated Flower Classification over a Large Number of Classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rakaraddi, A.; Siew-Kei, L.; Pratama, M.; and de Carvalho, M. 2024. Graph Mining under Data scarcity. *arXiv preprint arXiv:2406.04825*.
- Sagi, O.; and Rokach, L. 2018. Ensemble learning: A survey. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 8(4): e1249.
- Stallkamp, J.; Schlipsing, M.; Salmen, J.; and Igel, C. 2012. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32: 323–332.
- Wang, Q.; Farahat, A.; Gupta, C.; and Zheng, S. 2021. Deep time series models for scarce data. *Neurocomputing*, 456: 504–518.
- Wen, Y.; Tran, D.; and Ba, J. 2020. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. *arXiv preprint arXiv:2002.06715*.
- Wortsman, M.; Ilharco, G.; Gadre, S. Y.; Roelofs, R.; Gontijo-Lopes, R.; Morcos, A. S.; Namkoong, H.; Farhadi, A.; Carmon, Y.; Kornblith, S.; et al. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, 23965–23998. PMLR.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xiao, J.; Hays, J.; Ehinger, K. A.; Oliva, A.; and Torralba, A. 2010. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 3485–3492. IEEE.
- Xu, Z.; Yuan, K.; Wang, H.; Wang, Y.; Song, M.; and Song, J. 2024. Training-Free Pretrained Model Merging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5915–5925.
- Yadav, P.; Tam, D.; Choshen, L.; Raffel, C. A.; and Bansal, M. 2024. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36.
- Yang, E.; Wang, Z.; Shen, L.; Liu, S.; Guo, G.; Wang, X.; and Tao, D. 2023. Adamerging: Adaptive model merging for multi-task learning. *arXiv preprint arXiv:2310.02575*.
- Yu, L.; Yu, B.; Yu, H.; Huang, F.; and Li, Y. 2024. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*.
- Zhou, Z.; Yang, Y.; Yang, X.; Yan, J.; and Hu, W. 2023. Going beyond linear mode connectivity: The layerwise linear feature connectivity. *Advances in Neural Information Processing Systems*, 36: 60853–60877.