

# Weight Scope Alignment: A Frustratingly Easy Method for Model Merging

Yichu Xu<sup>ab</sup>, Xin-Chun Li<sup>ab</sup>, Le Gan<sup>ab</sup> and De-Chuan Zhan<sup>ab,\*</sup>

<sup>a</sup>School of Artificial Intelligence, Nanjing University, China

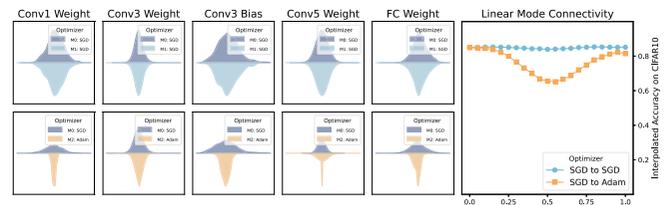
<sup>b</sup>National Key Laboratory for Novel Software Technology, Nanjing University, China

**Abstract.** Merging models becomes a fundamental procedure in some applications that consider model efficiency and robustness. The training randomness or Non-I.I.D. data poses a huge challenge for averaging-based model fusion. Previous research efforts focus on element-wise regularization or neural permutations to enhance model averaging while overlooking weight scope variations among models, which can significantly affect merging effectiveness. In this paper, we reveal variations in weight scope under different training conditions, shedding light on its influence on model merging. Fortunately, the parameters in each layer basically follow the Gaussian distribution, which inspires a novel and simple regularization approach named **Weight Scope Alignment (WSA)**. It contains two key components: 1) leveraging a target weight scope to guide the model training process for ensuring weight scope matching in the subsequent model merging. 2) fusing the weight scope of two or more models into a unified one for multi-stage model fusion. We extend the WSA regularization to two different scenarios, including Mode Connectivity and Federated Learning. Abundant experimental studies validate the effectiveness of our approach.

## 1 Introduction

As a technique for combining multiple deep models into a single model, model fusion [34, 24] [34] has gained widespread applications across various domains, including Mode Connectivity [4, 9, 10] and Federated Learning [1, 38, 51]. First, the model interpolation could shed light on the properties of the mode connectivity in neural networks [16, 11, 14]. Then, due to data privacy protection, transmitting intermediate models across edge nodes and fusing them on the server has been the common procedure in federated learning [46, 33, 37]. To be brief, model fusion matters a lot in these applications and has attracted a wide range of research interest. The primary goal of model fusion is to retain the capabilities of the original models while achieving improved generalization, efficiency, and robustness.

The coordinate-based parameter averaging is the most common approach for model fusion in deep neural networks [38, 34, 48]. The research on mode connectivity involves a linear or piece-wise interpolation between models [16, 45], while federated learning takes the averaging of local models from edge nodes for aggregation [38, 7]. Although parameter averaging exhibits favorable properties, it may not perform optimally in more complex training scenarios, especially when faced with various training conditions or Non-Independent



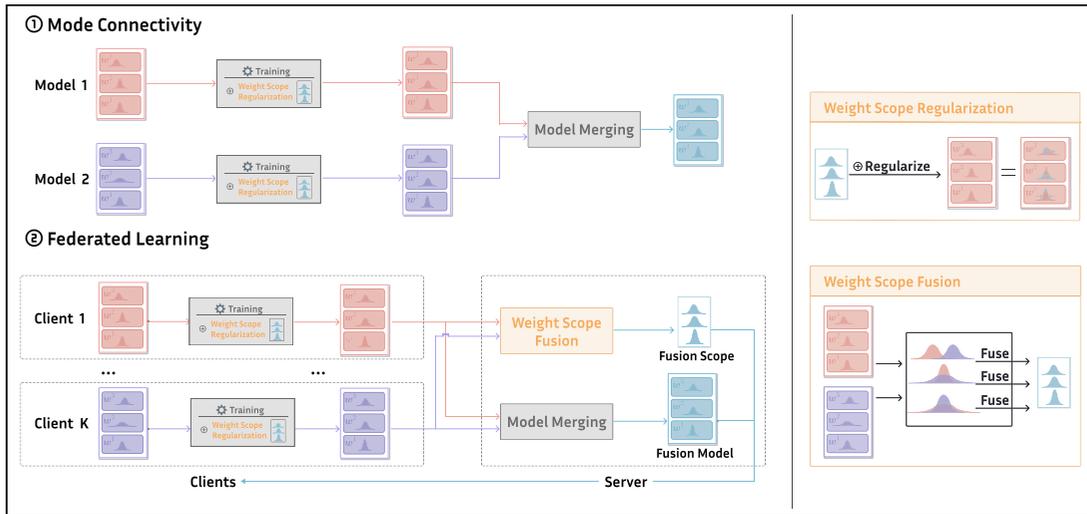
**Figure 1.** **Left:** Similar optimizers yield similar weight distributions. Conversely, different optimizers lead to distinct weight profiles. **Right:** Models trained from different training conditions (e.g., the used optimizer) tend to produce poorer model interpolations. Details are in Section 5.1.

and Identically Distributed (Non-I.I.D.) data. For example, the Non-I.I.D. data in federated learning means that the data of local nodes are naturally heterogeneous, making the model aggregation suffer from diverged update directions [20, 22]. Additionally, the property of permutation invariance that neural networks own exacerbates the challenge of model fusion because of the neuron misalignment phenomenon [10, 53, 5, 15]. Hence, solutions have been proposed from the aspect of element-wise regularization [33, 1, 7] or mitigating the permutation invariance [35, 3, 42, 40]. Few of these methods, however, have considered the impact of weight ranges across models on model fusion.

In this paper, we investigate the influence of different training conditions on model weight distributions (defined as Weight Scope) and further study model merging under various weight scopes. We first conduct several experiments under various training hyper-parameters or data quality conditions and find that the weight scopes of the converged models differ a lot, a phenomenon we define as “weight scope mismatch”. Figure 1 illustrates the model weight distributions under different training conditions, revealing noticeable differences despite all distributions being approximated by Gaussian distributions. Specifically, the top five sub-figures are parameters from models that use the same optimizer, while the bottom ones take different optimizers. In the rightmost of Figure 1, the linear interpolation results that reflect the mode connectivity property are also provided. Clearly, the mismatched weight scope leads to a worse linear interpolation, highlighting the impact of weight range inconsistency on model fusion. To intuitively explain, parameters with similar distributions can be aggregated more easily, whereas those with dissimilar distributions often present challenges in model merging.

Fortunately, the parameters in each layer follow a very simple distribution, i.e., Gaussian distribution. The simple distribution inspires

\* Corresponding Author. Email: zhandc@nju.edu.cn



**Figure 2.** The Framework of Weight Scope Alignment and Its Applications. **Left:** This method can be adapted to various applications of model merging, such as Model Connectivity and Federated Learning. Mode Connectivity includes a single model fusion, while Federated Learning requires multi-stage fusion. **Right:** The method comprises two components: weight scope regularization and weight scope fusion.

us a novel and easy way of aligning parameters. We leverage a target weight scope to guide the model training process, ensuring weight alignment and scope matching in subsequent model merging. For more complex multi-stage fusion, we calculate the mean and variance of parameter weights in the to-be-merged models, then aggregate these statistics into unified one as the target weight scope. The proposed method is named **Weight Scope Alignment (WSA)**, and the above steps are named weight scope regularization and weight scope fusion, respectively. The whole procedure is illustrated in Figure 2. We apply WSA to scenarios of mode connectivity and federated learning for exploring the advantages of WSA compared with corresponding solutions in these areas. Our work aims to align the weights as closely as possible with a given distribution during training, thereby enhancing the match in weight ranges and facilitating model merging. Compared with other similar regularization methods, e.g., the weight decay and the proximal term, the proposed WSA seeks to balance specificity and generality, addressing the shortcomings of existing methods while optimizing for effective model fusion. Our contributions are as follows: 1) *as far as we know, our work is the first to formally study the impact of weight scope on model fusion*; 2) *the proposed WSA is simple yet effective, which is verified on two applications via abundant experimental studies*.

## 2 Related Works

Model fusion is a fundamental technique in several applications. The related scenarios and solutions are introduced as follows.

**Model Merging in Mode Connectivity.** Visualizing loss landscape is an intuitive way to understand the mode connectivity [31, 17, 36], where the landscape is shown in a 2-dim or 3-dim space via model interpolations. The model interpolations between the initialization and the converged model could reflect the monotonic linear interpolation phenomenon [16, 11, 47, 43]. [13] bridges the model connectivity and lottery ticket hypothesis [12], proposing a method for model pruning. Mode connectivity is also related to the model optimization and generalization [27, 8].

The work [16] also points out that two independent minima suffer a barrier in their linear interpolation, which attracts several solu-

tions to decrease the barrier. [9, 14] make a notable discovery that the independent minima are possible to be connected via a simple piece-wise or quadratic curve. [39, 48] find that the minima fine-tuned from the same pre-trained model could mitigate the barrier in linear interpolation. [10] guesses that the independent minima are located in the same basin with the consideration of permutation invariance, conjecturing that the minima matched via the simulated annealing algorithm encounter no barrier. [42] propose the weight-based and activation-based matching method via the optimal transport. [44] decreases the barrier via both the permutation alignment and the quadratic curve. [3] employs three distinct neuron-matching methods to corroborate the low-barrier hypothesis. Although the permutation invariance is considered in these works, the mismatch of weight scope in neural networks is also fundamental to model fusion. Our proposed method could further decrease the barrier on the basis of these works.

**Model Merging in Federated Learning.** Federated learning (FL) aims to break the limitation of data privacy, utilizing a server that collaborates with client devices to train a model [51]. As the most standard algorithm in FL, FedAvg [38] takes a simple coordinate-based parameter averaging on the server to accomplish the model fusion process. A huge challenge that FL faces is the Non-I.I.D. data [20, 26], where the inherent data heterogeneity leads to weight divergence during local training [22]. Applying coordinate-wise regularization on local models is a popular solution to solve the Non-I.I.D. challenge in FL. [32] claims that weight decay [30] can lead to divergent optimization objectives among Non-I.I.D. clients in FL. [54, 33] introduce a proximal term in local optimization. This term helps align local optimization more closely with given model parameters, facilitating model aggregation. Additionally, [26] implements constraints on the local gradient directions of each client, nudging them closer to a global direction. However, the coordinate-wise regularization tends to be overly specific, which may induce perturbations to model training.

Additionally, the permutation invariance could also make the local models misaligned in FL, which is harder for aggregation. [53, 46] consider the permutation invariance of neural networks, rearranging the order of neurons can result in multiple models with the same

functionality as the original network. [52] introduces a group alignment method, while [35] designs the position-aware neurons to align parameters. These works mainly focus on the permutation invariance in FL, but they have not analyzed the influence of different model weight ranges on model fusion.

### 3 Proposed Methods

#### 3.1 Preliminaries

We introduce a collection of models, denoted as  $\mathcal{K}$ , with the size of the collection  $|\mathcal{K}| \geq 1$ . Each model within this collection is constructed based on the same underlying model architecture. The distinctiveness among the models in  $\mathcal{K}$  arises from the variations in their weight layers, which are derived from training under various hyperparameters or different data.

To encapsulate the statistical characteristics of the weight layers across different models, we denote the weight matrix of the  $\ell^{th}$  layer in the  $k^{th}$  model as  $\mathbf{w}_k^\ell$ . Our assumption is that the elements within this matrix follow a Gaussian distribution which is characterized by its mean  $\mu_k^\ell$  and standard deviation  $\sigma_k^\ell$ . The assumption is rational and we will empirically present the weight distributions of the converged model in Section 5.1. Formally, the distribution of the weight matrix  $\mathbf{w}_k^\ell$  is as follows:

$$p(\mathbf{w}_k^\ell) = \mathcal{N}(\mu_k^\ell, (\sigma_k^\ell)^2), \quad (1)$$

$$\mu_k^\ell = \frac{1}{|\mathbf{w}_k^\ell|} \sum_{w \in \mathbf{w}_k^\ell} w, \quad \sigma_k^\ell = \sqrt{\frac{1}{|\mathbf{w}_k^\ell|} \sum_{w \in \mathbf{w}_k^\ell} (w - \mu_k^\ell)^2}. \quad (2)$$

In Equation 2, we apply Maximum Likelihood Estimation to derive the standard deviation  $\sigma_k^\ell$  and mean  $\mu_k^\ell$  of the weight  $\mathbf{w}_k^\ell$ .

#### 3.2 Weight Scope Alignment

Firstly, we introduce Weight Scope Regularization. Then, for more complex multi-stage fusion, we propose Weight Scope Fusion for better target alignment.

**Weight Scope Regularization.** Given a weight distribution  $\mathcal{N}(\mu, \sigma^2)$  and a target weight distribution  $\mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)$ , our method endeavors to ensure consistency between them. This consistency is crucial for guaranteeing scope-matched distribution among the new models in the subsequent model fusion. To achieve this, we employ a divergence measure known as the Kullback-Leibler (KL) divergence, specifically focusing on calculating the divergence between the two univariate Gaussian distributions, which is given by:

$$D_{\text{KL}} = \log\left(\frac{\tilde{\sigma}}{\sigma}\right) + \frac{\sigma^2 + (\mu - \tilde{\mu})^2}{2\tilde{\sigma}^2} - \frac{1}{2}, \quad (3)$$

where  $\tilde{\mu}, \tilde{\sigma}$  are hyperparameters. By minimizing the KL divergence between the training models' weight distribution and their goal, it ensures that the weight distributions are closely aligned, facilitating more effective and harmonious model fusion.

**Weight Scope Fusion.** In some complex scenarios with large amounts of models and multiply stages of model merging, a given pre-defined weight distribution is not adaptable. Therefore, we propose a method named Weight Scope Fusion to enhance the applicability of Weight Scope Regularization. Focusing on the weights of the  $\ell^{th}$  layer, we assume that each weight  $\mathbf{w}_k^\ell$ ,  $k \in \mathcal{K}$ , follows its own Gaussian distribution, and they are independent of each other. We can get the fused Gaussian distribution  $\mathcal{N}(\tilde{\mu}^\ell, (\tilde{\sigma}^\ell)^2)$  by:

$$\tilde{\mu}^\ell = \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \mu_k^\ell, \quad \tilde{\sigma}^\ell = \frac{1}{\sqrt{|\mathcal{K}|}} \sqrt{\sum_{k \in \mathcal{K}} (\sigma_k^\ell)^2}. \quad (4)$$

#### 3.3 Analysis and Comparisons with Other Methods

We next provide some analysis of the proposed simple method, especially the comparisons with existing works.

**Comparison with Weight Decay.** Weight decay [30] stands as a cornerstone in model regularization, advocating for weight constraints that push the weights towards zero and promote uniformity. Considering a weight vector, denoted as  $\mathbf{w} \in \mathbb{R}^n$ , with elements  $[\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_n]$ . For ease of analysis, we define the mean, standard deviation, and L2 norm of this vector as follows:  $\mu = \frac{1}{n} \sum_{i=1}^n \mathbf{w}_i$ ,  $\sigma = \sqrt{\frac{1}{n} \sum_{i=1}^n (\mathbf{w}_i - \mu)^2}$ ,  $\|\mathbf{w}\|_2 = \sqrt{\sum_{i=1}^n \mathbf{w}_i^2}$ . The weight decay term is formulated as  $\frac{\lambda}{2} \|\mathbf{w}\|_2^2$ . We can express weight decay in terms of  $\mu$  and  $\sigma$  as follows:

$$\frac{\lambda}{2} \|\mathbf{w}\|_2^2 = \frac{\lambda n}{2} (\sigma^2 + \mu^2). \quad (5)$$

The weight decay term brings the mean and variance of model weights close to zero, and several research [1, 26] has designed FL algorithms with adjustment of weight decay term in local training. However, the impact of weight decay on models can not determine the shape of weight distributions and could not align parameter distributions under different conditions. Hence, it does not necessarily lead to a harmonization of the model scopes before fusion.

**Comparison with Proximal Term.** The proximal term is used to keep the model weight  $\mathbf{w}$  and another one  $\tilde{\mathbf{w}} \in \mathbb{R}^n$  as close as possible during the update of  $\mathbf{w}$ , which is represented as  $\|\mathbf{w} - \tilde{\mathbf{w}}\|_2^2$ . In FedProx [33], the proximal term constrains the local models to stay close to the global model, thereby ensuring stability during model fusion. For clarity, we also define  $\tilde{\mu}$  and  $\tilde{\sigma}$  as the mean and standard deviation of vectors  $\tilde{\mathbf{w}}$ . The proximal term can be decomposed as follows:

$$\|\mathbf{w} - \tilde{\mathbf{w}}\|_2^2 = n\sigma^2 + n\mu^2 + n\tilde{\sigma}^2 + n\tilde{\mu}^2 - 2 \sum_{i=1}^n \mathbf{w}_i \tilde{\mathbf{w}}_i. \quad (6)$$

Because the part of  $n\tilde{\sigma}^2 + n\tilde{\mu}^2$  can be seen as a constant term, the proximal term differs from weight decay by only an additional term of  $-2 \sum_{i=1}^n \mathbf{w}_i \tilde{\mathbf{w}}_i$ , which encourage the weights in  $\mathbf{w}$  to align with the direction of  $\tilde{\mathbf{w}}$  as closely as possible. However, the restriction is too strict because it requires the direction alignment for each specific element. This may limit the normal training process and perturb the effects of other loss functions.

In the context of model fusion, weight decay represents a more flexible approach, while the proximal term seems to introduce directional constraints additionally. Both of them have overlooked the influence of weight scope on model fusion and are unable to align weight scopes, as demonstrated in Figure 11. Our proposed method aims to achieve consistency in model weight ranges for improved model fusion results.

**Comparison with Network Invariance.** Some studies have investigated the impact of permutation invariance on mode connectivity [10, 42, 3, 44, 4, 15, 5] and model aggregation in federated learning [53, 35, 46]. A more detailed introduction of these studies has been presented in Section 2. However, the weight scope mismatch could make the alignment algorithm inaccurate, especially the weight-based ones [42]. Additionally, even if the neurons are aligned in order, their weight scopes are still mismatched, which could also

lead to performance degradation. Several works also study the scale invariance, e.g., [40] normalizes the model parameters layer-wisely and then searches the proper neural permutations, and [25] studies the relative scale of weight and bias in monotonic linear interpolation. None of these works formally studies the impact of weight scope on model fusion.

## 4 Applications of WSA

This section presents the applications of Weight Scope Alignment to mode connectivity and federated learning. They can be seen as one-stage and multi-stage model fusion scenarios, respectively. While the specific processes or the model set to be fused may differ across them, they all share the common thread of leveraging the proposed WSA as a constraint during training.

### 4.1 Mode Connectivity

Given two well-trained models,  $\mathbf{w}_1$  and  $\mathbf{w}_2$ , the loss barrier along their linear interpolation path is:

$$\max_{\alpha \in [0,1]} \mathcal{L}((1-\alpha)\mathbf{w}_1 + \alpha\mathbf{w}_2) - [(1-\alpha)\mathcal{L}(\mathbf{w}_1) + \alpha\mathcal{L}(\mathbf{w}_2)], \quad (7)$$

where  $\alpha$  is the interpolation coefficient, and the loss barrier represents the point of maximum loss increase along the linear interpolation path. A higher loss barrier suggests that the two models may not be in the same basin within the loss landscape, while a lower loss barrier indicates linear mode connectivity.

Both OTFusion [42] and Git Re-basin [3] have improved the way they perform model interpolation by considering neuron matching. They calculate matching relationships  $\Pi$  between the weights of each layer in the two models, using a permutation matrix or optimal transport matrix, leading to the fusion formula:

$$(1-\alpha)\mathbf{w}_1 + \alpha\Pi\mathbf{w}_2. \quad (8)$$

Our approach can easily be integrated into the aforementioned model interpolation methods to achieve improved mode connectivity. For separately trained models, they are typically randomly initialized from the same distribution, and their weight scopes are similar. The mean  $\tilde{\mu}$  and standard deviation  $\tilde{\sigma}$  of the target weight distribution are the hyperparameters and keep invariant during each model's training. We incorporate a Weight Scope Regularization term (i.e., Equation 3) to ensure that the weight range remains close to the target weight scope. With our method, the weight scope of the converged models are matched, which is beneficial for searching the matching matrix.

### 4.2 Federated Learning

Model fusion is a fundamental procedure in Federated Learning (FL), i.e., improving the performance of joint training by merging models trained on different data sources. Its challenge lies in the fact that real-world data distributions across different clients are heterogeneous, leading to a phenomenon known as "client drift" during local client training, which impacts the overall performance of FL. Some recent research has addressed this issue from a Non-I.I.D. perspective, proposing various improved methods.

We use the previously defined  $\mathcal{K}$  to represent the set of clients, with each client associated with the data distribution  $\mathcal{D}_k$ . The objective of FL is:

$$\min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w}), \quad f(\mathbf{w}) \triangleq \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \mathcal{L}(\mathbf{w}; \mathcal{D}_k), \quad (9)$$

---

### Algorithm 1 FedAvg with WSA

---

- 1: **Input:** model  $\mathbf{w}$ , number of rounds  $T$ , local iteration steps  $\tau$ , parameters  $\eta, \varepsilon$ ,
  - 2: **for**  $t = 0, \dots, T - 1$  communication rounds **do**
  - 3:   **Global server:**
  - 4:     Send  $\mathbf{w}, \tilde{\mathbf{p}} = \{\tilde{\mathbf{p}}^\ell = (\tilde{\mu}^\ell, \tilde{\sigma}^\ell)\}_{\ell \in [L]}$  to all clients
  - 5:   Client  $k \in \mathcal{K}$  in parallel **do:**
  - 6:     Set  $\mathbf{w}_k \leftarrow \mathbf{w}$
  - 7:     **for**  $s = 0, \dots, \tau - 1$  local iterations **do**
  - 8:       Update  $\mathbf{w}_k \leftarrow \mathbf{w}_k - \eta \nabla \mathcal{L}(\mathbf{w}_k; \mathcal{D}_k, \tilde{\mathbf{p}})$  using loss in Equation 10
  - 9:       Send  $\mathbf{w}_k, \{\mu_k^\ell, \sigma_k^\ell\}_{\ell \in [L]}$  to the server
  - 10:    **Global server:**
  - 11:     Update global model  $\mathbf{w} \leftarrow \frac{1}{|\mathcal{K}|} \sum_{k \in \mathcal{K}} \mathbf{w}_k$
  - 12:     Update  $\tilde{\mu}^\ell, \tilde{\sigma}^\ell$  using Equation 4
- 

where  $\mathcal{L}(\cdot)$  evaluates the loss for each data sample on model  $\mathbf{w}$ . Cross-entropy is commonly used as the loss function. To apply our method in FL, we additionally upload the weight scope (i.e., Equation 2) to the server after each local training procedure. The server then performs the fusion of these weight scopes (i.e., Equation 4) and sends the merged weight scopes  $\tilde{\mathbf{p}}$  back to the clients. During local client training, the optimization proceeds as follows:

$$\mathcal{L}_{\text{local}}(\mathbf{w}; \mathcal{D}_k, \tilde{\mathbf{p}}) = \mathcal{L}(\mathbf{w}; \mathcal{D}_k) + \lambda \sum_{\ell=1}^L D_{\text{KL}}(p(\mathbf{w}^\ell), \tilde{\mathbf{p}}^\ell), \quad (10)$$

where  $\lambda$  is used to control the strength of weight scope alignment, and  $\tilde{\mathbf{p}}^\ell$  is the fused weight scope of the  $\ell^{\text{th}}$  layer. The full framework of FedAvg [38] with WSA is outlined in Algorithm 1.

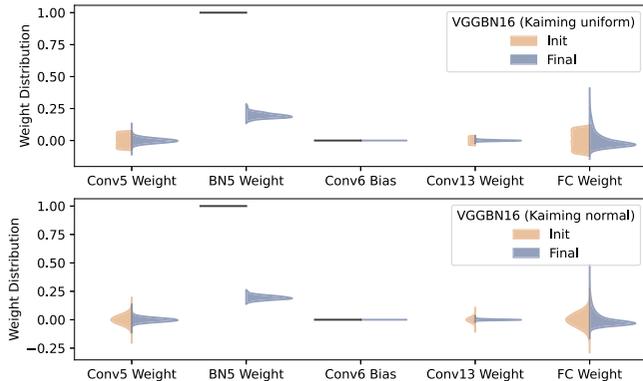
## 5 Experiments

In Section 5.1, we verify that different training conditions lead to variations in weight scope, and the differences between weight scope can affect the performance of model merging. In Section 5.2, we explore the effectiveness of the WSA method in the mode connectivity scenario. In Section 5.3, we demonstrate the performance improvements achieved by WSA in various federated learning scenarios and analyze the impact of the method.

### 5.1 Basic Experiments

**Observations.** Our basic assumption is that the weight scope could be formulated as the Gaussian distribution. Hence, we first study the weight distributions of the models. The first observation is that the converged weight scope is irrelevant to the way of weight initialization, e.g., the Kaiming uniform or the Kaiming normal initialization method [18]. Figure 3 shows the weight distributions of several layers in VGG16 [41] with BatchNorm [21], where we train the network on CIFAR-100 [29] for 200 epochs. In the supplementary material [49], we present the weight scopes in different training settings, which also follow the Gaussian distributions.

**The Influence of Weight Scope on Model Fusion.** We then investigate the difference of weight scope under different conditions, which include: 1) hyperparameters, e.g., optimizer, batch-size, learning rate, and weight decay; 2) data quality, e.g., label imbalance, label noise, feature noise, and data size. For each condition, we select two specific settings, and then train models under the same setting or not. For example, the optimizer could be SGD or Adam [28], and we



**Figure 3.** The Gaussian distribution of parameters in VGGBN16 trained on CIFAR-100. Both uniform and normal initialization lead to Gaussian parameter distributions.

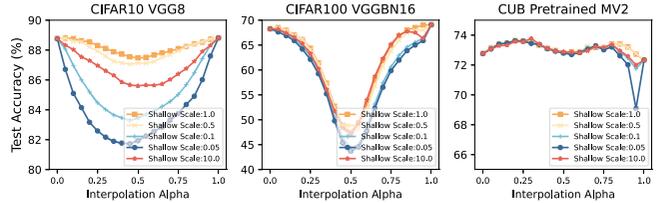
**Table 1.** The impact of training conditions on LMC.

Factor	KL D. ( $\times 10^{-4}$ )	Ba(=) (%)	Ba( $\neq$ ) (%)	Diff (%)
Label Imb.	9.60	4.60	22.05	17.45
Optimizer	8.95	1.13	17.84	16.71
Batch Size	5.96	44.98	56.07	11.09
Label Noise	1.49	0.58	0.79	0.21
Learning Rate	0.67	0.38	6.77	6.39
Feature Noise	0.60	0.70	9.60	8.90
Data Size	0.37	0.00	0.75	0.75
Weight Decay	0.31	1.59	2.31	0.72

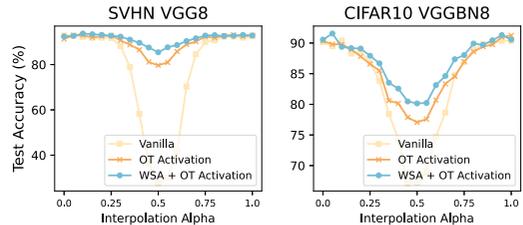
train three models with each using SGD, SGD, and Adam as the optimizer, respectively. Then, we plot the weight scopes of these models and investigate the linear interpolation accuracy. Figure 1 shows that models trained using different optimizers own varying weight scopes and the interpolation meets an obvious barrier. The details of training conditions and the corresponding illustration results can be found in the supplementary material [49].

From Figure 1, we can observe that the weight scopes under the same training condition are near the same. In fact, we calculate their KL divergence, and the results are near zero. Hence, we calculate the average KL divergence of weight scopes under different conditions, e.g., the average divergence of Gaussian distributions in the bottom five sub-figures (the “KL D.” column in Table 1). Additionally, we calculate the barriers of the two interpolated curves, and their difference, which are listed in columns of “Ba(=)”, “Ba( $\neq$ )”, and “Diff” in Table 1. In the table, the “Ba(=)” column is commonly smaller than “Ba( $\neq$ )”, which verifies that models under the same condition are indeed similar in weight scopes. Empirically, for models under different conditions, a larger KL divergence corresponds to a larger barrier. This shows that the weight scope mismatch is strongly correlated with the model fusion performance.

**Model Fusion under Manual Scaling.** The above experiments could not support that the weight scope is the cause of the barrier in model fusion. Thanks to the scale invariance of neural networks, we could manually scale the networks and create the weight scope mismatch phenomenon. Specifically, we train two networks under the same condition with different random initialization. To avoid overuse of symbols, we denote them as  $w_1$  and  $w_2$ , respectively. To create the scope mismatch, we select one layer in  $w_2$  and multiply its weight by  $\alpha$ , and divide its following layer’s weight by  $\alpha$ . The obtained model



**Figure 4.** The linear interpolation curves between a model with another model (Scale=1.0) and its scaled versions (Scale  $\neq$  1.0).



**Figure 5.** Model interpolation on SVHN and CIFAR-10. WSA enhances OTFusion [42] performance.

is denoted as  $w_2 \cdot \alpha$ , which performs the same as  $w_2$ . However, when taking interpolations between  $w_1$  and  $w_2 \cdot \alpha$ , the interpolation curves differ a lot. The results are shown in Figure 4, where the “Scale” in the legend denotes  $\alpha$ . Using  $\alpha = 1.0$  means no scaling. Clearly, scaling the layers could make the barrier more obvious, especially in VGG8. Experimental details and more analysis can be found in the supplementary material [49].

## 5.2 Performance in Mode Connectivity

To investigate the effectiveness of our approach on mode connectivity, we apply the proposed WSA to OTFusion [42] and Git Re-Basin [3]. Specifically, we first initialize and train two models, and plot the vanilla interpolation curve. Then, we use the activation-based OTFusion method to search for an alignment and plot the interpolated curve after matching. Finally, we replace the models with another two models trained using the WSA as introduced in Section 4.1 and also use the OTFusion to search the alignment matrix. The results are shown in Figure 5. OTFusion could decrease the barrier of vanilla interpolation, and our proposed WSA could improve its performance further.

Then, we combine our WSA with Git Re-Basin, which searches the permutation matrix instead of an optimal transport matrix compared with OTFusion. Similar to OTFusion, we also train models with or without our proposed WSA. In Figure 6, we conduct experiments on a ResNet-22x2 [19, 3] model on CIFAR-10 [29] to compare the loss barrier at different epochs. The solid lines represent the vanilla model interpolation, while the dashed lines represent the interpolation after using Git Re-basin. On the one hand, applying Git Re-Basin could indeed decrease the barrier of vanilla interpolation. However, as the number of epochs increases, the results of w/o WSA continue to grow, even when permutation is considered. In contrast, the lines of w/ WSA can consistently maintain lower loss barriers. In the middle and right figures, we compare loss barriers under different model widths and depths. Specifically, we compare ResNet-22 models with a widening factor of 2, 4, and 6, and ResNet models with a widening factor of 2 and depths of 16, 22, and 28. The results show that our method consistently improves model fusion across models

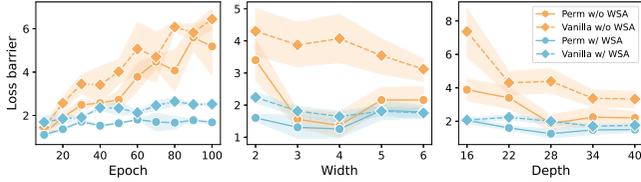


Figure 6. WSA facilitates Git Re-Basin [3] under different epochs, widths and depths, ensuring a smaller loss barrier.

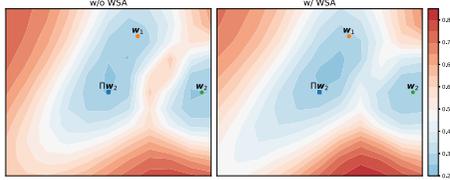


Figure 7. Loss landscape w/o WSA and w/ WSA.

with varying depths and widths. Furthermore, permutation interpolation further improves the loss barrier, demonstrating that considering both weight permutation and scope can effectively reduce the loss barrier. This indicates that WSA helps align the weight ranges of models, contributing to enhanced mode connectivity.

In Figure 7, we illustrate the loss landscape of model  $w_1$ ,  $w_2$  and  $\Pi w_2$ , where the permutation matrix  $\Pi$  is computed by Git Re-basin. We train the two models with and without our proposed WSA, separately. First, there is an obvious peak between  $w_1$  and  $w_2$  on the left; meanwhile, on the right, the loss between  $w_1$  and  $w_2$  is alleviated with WSA. Secondly, through Git Re-basin, the valley between  $\Pi w_2$  and  $w_1$  is a larger flat area which leads to better model merging.

**Pretrained Model.** We conduct experiments using the Vision Transformer (ViT) pre-trained on ImageNet-21k on mode connectivity. Figure 8 presents the results of finetuning the pre-trained ViT on ImageNet-100. Specifically, We finetune the ViT using two distinct random seeds and then conduct model interpolation between the resultant models. “baseline” refers to training without intervention, while “w/ WSA” indicates training with Weight Scope Alignment. Because of permutation-based weight averaging methods are not designed for transformers, we employ the naive model interpolation  $(1 - \alpha)w_1 + \alpha w_2$ . The results demonstrate that WSA can improve the performance of interpolated models.

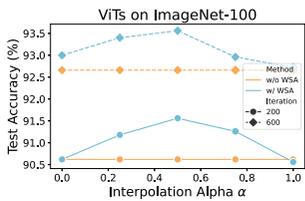


Figure 8. Mode Connectivity of ViTs finetuned on ImageNet-100.

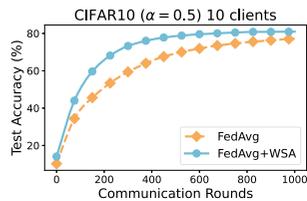


Figure 9. Test accuracy curve of VGG8 trained on CIFAR-10 in FL.

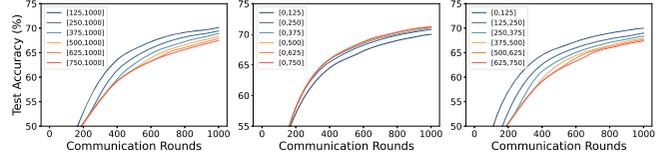


Figure 10. The Impact of using WSA in different rounds in FL.

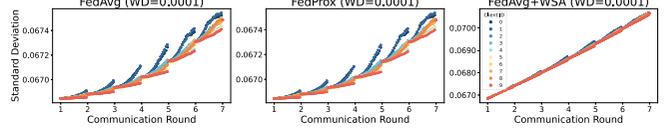


Figure 11. The standard deviation of conv2 weights for each client in FL training with 10 clients, varying across rounds.

### 5.3 Performance in Federated Learning

To simulate real-world federated learning (FL) scenarios, we initially consider image classification tasks across three datasets: CIFAR-10 [29], CIFAR-100 [29], and CINIC-10 [6]. We explore two different client scenarios: “cross-device” with 100 small-data clients where only 10 clients participate in each training round, and “cross-silo” with 10 clients, each having larger datasets and engaging each training round. We use the same CNN architecture used in [26, 23].

To ensure Non-I.I.D. data distribution in data partitioning, we employ the Dirichlet distribution for creating heterogeneous data on each client [52, 35]. We compare with several state-of-the-art (SOTA) methods, including FedAvg [38], FedProx [33], SCAFFOLD [26], and FedExp [23]. Throughout all experiments, unless otherwise specified, we use a default batch size of 50, 20 local update steps, and 1500 communication rounds. Additional details and results are provided in the supplementary material [49].

Table 2 presents the performance comparison of CNN model trained in various FL settings with 10 clients where FedAvg+WSA outperforms FedAvg in performance. Moreover, as a plug-in, we can easily adapt it to existing FL methods, and the experiments indicate that incorporating WSA leads to significant improvements in each method. This underscores the effectiveness of our approach in model fusion of FL. Table 3 shows that similar results are observed with 100 clients. Furthermore, Figure 9 illustrates the adaptability of WSA to deeper models on CIFAR-10.

To further understand the impact of WSA on model merging, we conduct experiments on the CIFAR-10 dataset over various durations. We consider three different intervention times during training: from the start of training until iteration  $[0, t_1]$ , from iteration  $[t_1, T]$  until convergence, and during the iterations within the interval  $[t_1, t_2]$ . Figure 10 displays the convergence curves for these three scenarios. In the first scenario, we observe that the longer the intervention time, the better our method enhances the convergence performance. In the second one, we find that employing our method earlier results in greater performance improvements and faster convergence. Lastly, even if our method only intervenes during the initial  $[0, 125]$  rounds, it still significantly boosts performance. Drawing inspiration from the concept of a “critical learning period” [50, 2], we conclude that assisting weight alignment in the early stages of learning provides more substantial benefits for the model fusion scenario.

**Table 2.** Comparison of Test Accuracy (%) across various FL settings on three datasets, involving 10 clients: “Baseline” represents the original method, while “+WSA” indicates the incorporation of WSA into “Baseline”. Alpha ↓ denotes client heterogeneity ↑. **Bold** fonts highlight the optimal results between “Baseline” and “+WSA”. The results demonstrate that +WSA enhances performance in all FL settings.

Dataset	CIFAR-10				CIFAR-100				CINIC-10			
	0.5		1		0.5		1		0.5		1	
Algorithm	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA
FedAvg [38]	68.42	<b>72.39</b>	69.06	<b>73.09</b>	30.66	<b>35.23</b>	30.86	<b>36.16</b>	53.05	<b>57.88</b>	53.94	<b>58.53</b>
FedProx [33]	68.40	<b>72.49</b>	69.11	<b>73.18</b>	30.47	<b>35.39</b>	31.20	<b>36.24</b>	52.57	<b>57.54</b>	53.34	<b>58.26</b>
SCAFFOLD [26]	66.69	<b>71.61</b>	68.14	<b>72.65</b>	29.58	<b>34.38</b>	31.12	<b>36.76</b>	50.88	<b>55.87</b>	52.13	<b>57.52</b>
FedExp [23]	73.24	<b>75.55</b>	72.73	<b>76.10</b>	38.93	<b>41.71</b>	40.48	<b>43.16</b>	56.62	<b>60.77</b>	56.30	<b>59.93</b>

**Table 3.** Comparison of Test Accuracy (%) across various FL settings on three datasets, involving 100 clients where 10% participate in each round: “Baseline” represents the original method, while “+WSA” indicates the incorporation of WSA into “Baseline”. Alpha ↓ denotes client heterogeneity ↑. **Bold** fonts highlight the optimal results between “Baseline” and “+WSA”. The results demonstrate that +WSA enhances performance in all FL settings.

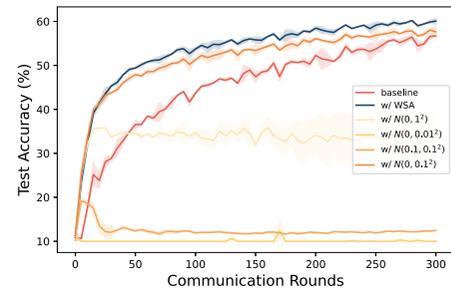
Dataset	CIFAR-10				CIFAR-100				CINIC-10			
	0.5		1		0.5		1		0.5		1	
Algorithm	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA	Baseline	+WSA
FedAvg	66.77	<b>71.11</b>	69.06	<b>72.80</b>	29.57	<b>34.47</b>	30.45	<b>35.31</b>	50.85	<b>56.68</b>	52.64	<b>57.82</b>
FedProx	66.82	<b>71.14</b>	69.06	<b>72.73</b>	29.72	<b>34.41</b>	30.35	<b>35.17</b>	50.14	<b>56.31</b>	52.13	<b>57.55</b>
SCAFFOLD	65.74	<b>70.36</b>	69.47	<b>73.25</b>	29.37	<b>34.85</b>	31.81	<b>37.31</b>	48.39	<b>54.49</b>	52.84	<b>57.07</b>
FedExp	71.79	<b>74.33</b>	71.82	<b>74.80</b>	37.98	<b>41.34</b>	38.46	<b>41.58</b>	50.22	<b>56.99</b>	55.84	<b>60.55</b>

**Analysis of weight distribution.** In Figure 11, we illustrate the alignment effect introduced by WSA on parameter ranges during the training process of FedAvg. The figure displays how the weight distributions of 10 clients evolve during training. To enhance clarity, we represent the distributions separately using mean and variance, showing the results for the original FedAvg, FedProx, and FedAvg+WSA. Different colors represent different clients, and aggregation occurs when the server averages the model parameters across all clients after each communication round, making the model distributions identical at that point. In the variance plots, we observe aggregation and divergence in FedAvg even with a weight decay of 0.0001, particularly in the early stages, indicating that the mismatch of weight ranges may contribute to slow convergence in the early phases of FL. In contrast, our constraint promotes a more cohesive model variance.

**Comparison with pre-defined distributions.** To demonstrate the effectiveness of weight scope fusion, we try to design some pre-defined distributions to replace it. Given that Gaussian distributions match the frequently noted attributes of model weights, we devise various pre-defined Gaussian distributions  $N(\mu, \sigma^2)$ , by altering the mean  $\mu$  and standard deviation  $\sigma$ . Each of these distributions is utilized to substitute the fusion distributions in WSA, and are set across all layers. As shown in Figure 12, our method “w/ WSA” outperforms all other comparative approaches. The pre-defined distribution  $w/ N(0, 0.1^2)$  is the second-best result, while other pre-defined ones led to a degradation in performance. We believe that a well-designed fix distribution can be beneficial during the initial training stage. Yet, weight distribution variation significantly shifts throughout the training, as seen in Figure 11, necessitating a dynamic and adaptable fusion distribution. Hence, our suggested weight scope fusion becomes crucial for broad model fusion applications.

## 6 Conclusion

In this study, we investigate the impact of different weight ranges of model parameters on model merging. We observe that models trained under various conditions exhibit inconsistencies in their



**Figure 12.** Comparison with pre-defined distributions in FL.

weight ranges, leading to a decrease in fusion performance. To address this issue, we propose Weight Scope Alignment method, which utilize Weight Scope Regularization to constrain the alignment of weight scopes during training, ensuring that the model’s weights closely match the specified scope. Moreover, for multi-stage fusion scenarios, we design Weight Scope Fusion to fuse the weight scopes of multiple models and regards the fused one as the target weight scope for Weight Scope Regularization. This alignment enhances performance during model merging. We validate the effectiveness of our approach in mode connectivity and federated learning.

## Acknowledgements

This research was supported by National Science and Technology Major Project (2022ZD0114805), NSFC (61773198, 62376118, 61921006), Collaborative Innovation Center of Novel Software Technology and Industrialization.

## References

- [1] D. A. E. Acar, Y. Zhao, R. M. Navarro, M. Mattina, P. N. Whatmough, and V. Saligrama. Federated learning based on dynamic regularization. In *ICLR*, 2021.
- [2] A. Achille, M. Rovere, and S. Soatto. Critical learning periods in deep networks. In *ICLR*, 2019.

- [3] S. K. Ainsworth, J. Hayase, and S. S. Srinivasa. Git re-basin: Merging models modulo permutation symmetries. In *ICLR*, 2023.
- [4] S. C. Ashmore and M. S. Gashler. A method for finding similarity between multi-layer perceptrons by forward bipartite alignment. In *IJCNN*, pages 1–7, 2015.
- [5] J. Brea, B. Simsek, B. Illing, and W. Gerstner. Weight-space symmetry in deep networks gives rise to permutation saddles, connected by equal-loss valleys across the loss landscape. *CoRR*, abs/1907.02911, 2019.
- [6] L. N. Darlow, E. J. Crowley, A. Antoniou, and A. J. Storkey. Cifar-10 is not imagenet or cifar-10. *CoRR*, abs/1810.03505, 2018.
- [7] C. T. Dinh, N. H. Tran, and T. D. Nguyen. Personalized federated learning with moreau envelopes. In *NeurIPS*, 2020.
- [8] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio. Sharp minima can generalize for deep nets. In *ICML*, volume 70, pages 1019–1028, 2017.
- [9] F. Draxler, K. Veschgini, M. Salmhofer, and F. A. Hamprecht. Essentially no barriers in neural network energy landscape. In *ICML*, volume 80, pages 1308–1317, 2018.
- [10] R. Entezari, H. Sedghi, O. Saukh, and B. Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. In *ICLR*, 2022.
- [11] J. Frankle. Revisiting "qualitatively characterizing neural network optimization problems". *CoRR*, abs/2012.06898, 2020.
- [12] J. Frankle and M. Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, 2019.
- [13] J. Frankle, G. K. Dziugaite, D. M. Roy, and M. Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *ICML*, volume 119, pages 3259–3269, 2020.
- [14] T. Garipov, P. Izmailov, D. Podoprikin, D. P. Vetrov, and A. G. Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *NeurIPS*, pages 8803–8812, 2018.
- [15] C. Godfrey, D. Brown, T. Emerson, and H. Kvinge. On the symmetries of deep learning models and their internal representations. In *NeurIPS*, 2022.
- [16] I. J. Goodfellow and O. Vinyals. Qualitatively characterizing neural network optimization problems. In *ICLR*, 2015.
- [17] Y. Hao, L. Dong, F. Wei, and K. Xu. Visualizing and understanding the effectiveness of bert. In *EMNLP-IJCNLP*, pages 4141–4150, 2019.
- [18] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, pages 1026–1034, 2015.
- [19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [20] K. Hsieh, A. Phanishayee, O. Mutlu, and P. B. Gibbons. The non-iid data quagmire of decentralized machine learning. In *ICML*, volume 119, pages 4387–4398, 2020.
- [21] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, volume 37, pages 448–456, 2015.
- [22] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. *CoRR*, abs/1811.11479, 2018.
- [23] D. Jhunjhunwala, S. Wang, and G. Joshi. Fedexp: Speeding up federated averaging via extrapolation. In *ICLR*, 2023.
- [24] X. Jin, X. Ren, D. Preotiu-Pietro, and P. Cheng. Dataless knowledge fusion by merging weights of language models. In *ICLR*, 2023.
- [25] K. Jordan, H. Sedghi, O. Saukh, R. Entezari, and B. Neyshabur. Repair: Renormalizing permuted activations for interpolation repair. In *ICLR*, 2023.
- [26] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *ICML*, volume 119, pages 5132–5143, 2020.
- [27] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *ICLR*, 2017.
- [28] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [29] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [30] A. Krogh and J. A. Hertz. A simple weight decay can improve generalization. In *NeurIPS*, pages 950–957, 1991.
- [31] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the loss landscape of neural nets. In *NeurIPS*, pages 6391–6401, 2018.
- [32] J. Li, A. Li, C. Tian, Q. Ho, E. P. Xing, and H. Wang. Fednar: Federated optimization with normalized annealing regularization. *CoRR*, abs/2310.03163, 2023.
- [33] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith. Federated optimization in heterogeneous networks. In *MLSys*, 2020.
- [34] W. Li, Y. Peng, M. Zhang, L. Ding, H. Hu, and L. Shen. Deep model fusion: A survey. *CoRR*, abs/2309.15698, 2023.
- [35] X.-C. Li, Y. Xu, S. Song, B. Li, Y. Li, Y. Shao, and D.-C. Zhan. Federated learning with position-aware neurons. In *CVPR*, pages 10072–10081, 2022.
- [36] X.-C. Li, J. Tang, B. Zhang, L. Li, and D.-C. Zhan. Exploring and exploiting the asymmetric valley of deep neural networks. *ArXiv*, abs/2405.12489, 2024.
- [37] Z. Li, H.-Y. Chen, H. W. Shen, and W.-L. Chao. Understanding federated learning through loss landscape visualizations: A pilot study. In *NeurIPS*, 2022.
- [38] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, volume 54, pages 1273–1282, 2017.
- [39] B. Neyshabur, H. Sedghi, and C. Zhang. What is being transferred in transfer learning? In *NeurIPS*, 2020.
- [40] F. Pittorino, A. Ferraro, G. Perugini, C. Feinauer, C. Baldassi, and R. Zecchina. Deep networks on toroids: Removing symmetries reveals the structure of flat regions in the landscape geometry. In *ICML*, volume 162, pages 17759–17781, 2022.
- [41] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [42] S. P. Singh and M. Jaggi. Model fusion via optimal transport. In *NeurIPS*, 2020.
- [43] B. Tao, X.-C. Li, and D.-C. Zhan. MLI formula: A nearly scale-invariant solution with noise perturbation. In *ICML*, 2024.
- [44] N. J. Tatro, P.-Y. Chen, P. Das, I. Melnyk, P. Sattigeri, and R. Lai. Optimizing mode connectivity via neuron alignment. In *NeurIPS*, 2020.
- [45] T. J. Vlaar and J. Frankle. What can linear interpolation of neural network loss landscapes tell us? In *ICML*, volume 162, pages 22325–22341, 2022.
- [46] H. Wang, M. Yurochkin, Y. Sun, D. S. Papailiopoulos, and Y. Khazaeni. Federated learning with matched averaging. In *ICLR*, 2020.
- [47] X. Wang, A. N. Wang, M. Zhou, and R. Ge. Plateau in monotonic linear interpolation - a "biased" view of loss landscape for deep networks. In *ICLR*, 2023.
- [48] M. Wortsman, G. Ilharco, S. Y. Gadre, R. Roelofs, R. G. Lopes, A. S. Morcos, H. Namkoong, A. Farhadi, Y. Carmon, S. Kornblith, and L. Schmidt. Model soups: Averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, volume 162, pages 23965–23998, 2022.
- [49] Y. Xu, X.-C. Li, L. Gan, and D.-C. Zhan. Weight scope alignment: A frustratingly easy method for model merging. 2024. URL <https://arxiv.org/abs/2408.12237>. Full version of this paper.
- [50] G. Yan, H. Wang, and J. Li. Critical learning periods in federated learning. *CoRR*, abs/2109.05613, 2021.
- [51] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.*, 10(2): 12:1–12:19, 2019.
- [52] F. Yu, W. Zhang, Z. Qin, Z. Xu, D. Wang, C. Liu, Z. Tian, and X. Chen. Fed2: Feature-aligned federated learning. In *KDD*, pages 2066–2074, 2021.
- [53] M. Yurochkin, M. Agarwal, S. Ghosh, K. H. Greenewald, T. N. Hoang, and Y. Khazaeni. Bayesian nonparametric federated learning of neural networks. In *ICML*, volume 97, pages 7252–7261, 2019.
- [54] S. Zhang, A. Choromanska, and Y. LeCun. Deep learning with elastic averaging sgd. In *NeurIPS*, pages 685–693, 2015.