

# Training-Free Pretrained Model Merging

Zhengqi Xu<sup>1</sup>, Ke Yuan<sup>1</sup>, Huiqiong Wang<sup>2</sup>, Yong Wang<sup>3</sup>, Mingli Song<sup>1</sup>, Jie Song<sup>1\*</sup>  
<sup>1</sup>Zhejiang University

<sup>2</sup>Ningbo Innovation Center, Zhejiang University

<sup>3</sup>State Grid Shandong Electric Power Company

{xuzhengqi, keyuan, huiqiong\_wang, brooksong, sjie}@zju.edu.cn, wangyong@sd.sgcc.com.cn

## Abstract

Recently, model merging techniques have surfaced as a solution to combine multiple single-talent models into a single multi-talent model. However, previous endeavors in this field have either necessitated additional training or fine-tuning processes, or require that the models possess the same pre-trained initialization. In this work, we identify a common drawback in prior works w.r.t. the inconsistency of unit similarity in the weight space and the activation space. To address this inconsistency, we propose an innovative model merging framework, coined as merging under dual-space constraints (MuDSC). Specifically, instead of solely maximizing the objective of a single space, we advocate for the exploration of permutation matrices situated in a region with a unified high similarity in the dual space, achieved through the linear combination of activation and weight similarity matrices. In order to enhance usability, we have also incorporated adaptations for group structure, including Multi-Head Attention and Group Normalization. Comprehensive experimental comparisons demonstrate that MuDSC can significantly boost the performance of merged models with various task combinations and architectures. Furthermore, the visualization of the merged model within the multi-task loss landscape reveals that MuDSC enables the merged model to reside in the overlapping segment, featuring a unified lower loss for each task. Our code is publicly available at [https://github.com/zju-vipa/training\\_free\\_model\\_merging](https://github.com/zju-vipa/training_free_model_merging).

## 1. Introduction

Great success has been achieved on various challenging tasks in computer vision by using deep neural networks, and a plethora of deep neural networks are developed and released publicly, with either their architectures and trained parameters (e.g., Pytorch Hub<sup>1</sup>, Hugging Hub<sup>2</sup>). These off-the-shelf

\*Corresponding author.

<sup>1</sup><https://pytorch.org/hub/>

<sup>2</sup><https://huggingface.co/HUB>




Figure 1. We use an intermediate layer of ResNet50 [16] to construct two groups of units as two parents, and then merge them pairwise. In the figure, each point represents a merged unit. The x-axis and the y-axis denote the weight and activation similarity between their two parents respectively. The color indicates the smaller activation similarity between them and their two parents. The star depicts the best merged unit which is most activation-similar to its parents. More details are provided in the supplementary material.

models are finely-tuned for a wide range of tasks, providing users with substantial convenience. However, these models are restricted to the tasks they were trained on, and this limitation poses significant challenges in terms of model storage [11, 63] and computation as the parameters of models continue to grow rapidly.

Given the abundance of well-trained models across various tasks, numerous research studies propose the amalgamation of multiple models into a single model that possesses multiple functionalities simultaneously. The existing body of literature can be broadly classified into two schools: *training-based knowledge amalgamation* [7, 14, 20, 30, 31, 37, 45,

46, 54, 56] and *training-free model merging* [18, 52, 53]. The former is inspired by knowledge distillation and adopts the outputs from multiple pretrained teachers to train a student model, which amalgamates the knowledge from multiple models into a single one. However, these methods usually require an additional expensive training process to facilitate the transfer of knowledge from the teacher models to the student model. In contrast, The latter fully utilize the parameters of the pretrained models to construct the multi-talent model, thereby avoiding additional training. The mainstream involves constructing multi-task models by adding task vectors, which are computed by subtracting the pre-trained model’s parameter values from those of the fine-tuned model, together. However, these methods can only merge fine-tuned models derived from the same pre-trained initialization, that is an extremely limited prerequisites. Emerging work *w.r.t.* unit matching [1, 22, 40, 44] shows promise to breaking the extra-training and pre-training curse, which has already reached zero-loss barriers in single-task model merging [1]. However, the permutation invariance [10], a cornerstone for most unit matching methods, seems not applicable to merge models from multiple tasks and leads to a poor artifact.

In this work, we have cast doubt upon the matched units identified by existing unit matching algorithms. Reviewing prior work, current unit matching methods either represent units using weights or activation values [27]. Nonetheless, using either weights or activation values as the matching criterion is incomplete. In terms of weights, the weight of a unit reflects its connection strength with the units in the previous layer, but it cannot capture the information of specific features (*i.e.* scale, offset) due to the absence of input data. As shown in Fig. 1, there are many pairs of units with highly similar weights but entirely different activation values, leading to a significant change in the function of their merged units. In terms of activations, it is compelling to measure units by their activations since two models must learn similar features to accomplish the same task [28], yet the same task can be achieved by distinct parameters [12]. Moreover, a phenomenon is shown in Fig. 1 that the best-merged unit is one whose parents simultaneously exhibits high similarity in both weight and activation.

Motivated by this, we propose a concise and effective model merging framework, termed *merging under dual-space constraints* (*MuDSC*) to achieve a more precise matching of relevant units. Specifically, the proposed *MuDSC* constrains the matching algorithm to find more suitable matches by linearly combining the similarity matrix of representations of units in both activation space and weight space. For the sake of rigor, we provide a formula derivation to prove the equivalence of this approach with a linear combination of optimization objectives in separate spaces. Moreover, we devise a novel iterative algorithm that simultaneously maximize the global similarity of matched units in both

weight space and activation space. In addition, we devise two variant of matching algorithms proposed in [1] and [43] to support the use of unit matching in networks with a group structure *i.e.* ViT [9] and Group Normalization [50].

In summary, the key contributions of this paper are:

1. We highlight the inconsistency of unit similarity in weight space and activation space, demonstrating through a straightforward experiment the impact of this inconsistency on model merging.
2. To take into account this inconsistency, we propose a novel model merging framework *MuDSC* (Merging under Dual-Space Constraints) that encourages both the activation spaces and weight spaces between two models be of high similarity.
3. We present a vigorous experimental study and show that *MuDSC* can significantly boost the merged performance in manifold multi-task scenarios for various architectures. Further loss landscape visualization verifies that our solution locates in the overlapping part with unified lower loss of each task.

## 2. Related Work

**Training-based model merging.** In order to reuse trained models and reduce the cost of training the new models from scratch, Shen *et al.* [37] propose knowledge amalgamation (KA), the goal of which is to learn a versatile student model from multiple task-specific teachers, but it is only verified on the comprehensive classification tasks. Following the knowledge amalgamation, Ye *et al.* [56] validate the effectiveness of KA in complex scenarios such as deep estimation and scene analysis, and in another work [57], they customized student models by amalgamating filtered knowledge from different teachers, even surpassing the those teachers in certain cases. There are more learning strategies have been proposed to enhance the performance [30, 31, 38] and usability [14, 20, 21, 58, 61] of KA. Although the success achieved by KA, a completely re-parameterized student model evidently leads to inefficient model merging. In contrast, model stitching [26] more efficiently leverages pre-trained models. model stitching aims to “plug-in” the bottom layers of one network into the top layers of another network, thus forming a stitched network [2, 6]. For instance, Yang *et al.* [55] introduce a novel two-stage strategy for reassembling customized networks from a zoo of pre-trained models under user-specified constraints, and Pan *et al.* [34] propose Stitchable Neural Networks for elastic deep learning by directly utilising the pretrained model families in model zoo via model stitching. Regrettably, the fine-tuning process is still necessary for bridging the gaps between pre-trained blocks in model stitching. Unlike KA and model stitching, we do not involve any additional training or fine-tuning processes in our merging methods.

**Training-free model merging.** Considering the data privacy

and the costly retraining for expanding model capabilities, recent efforts have sparked significant interest in exploring how to merge models from multiple tasks into the single multi-task model without additional training [18, 19]. Ilharco *et al.* [18] present a paradigm for editing neural networks, namely Task Arithmetic, that constructs multi-task models with performing simple addition operations on task vectors. Expanding on this groundwork, a series of work expand the scope of the task arithmetic framework [62] or solve the potential issues of task vector-based method [52, 53]. However, task vector-based methods can only be applied to models fine-tuned from the same pre-trained initialization, and this assumption severely restricts the applicability of model merging. Fortunately, unit matching technologies show promise in relaxing this assumption [1]. Guided by the existence of the permutation symmetries of hidden units in neural networks [4, 10, 13, 15, 17, 39, 44], unit matching aims at mapping all models to the same basin by permutating neuron units. Though we find it widely applied for merging models in the single-task scenarios [1, 22, 40, 48, 59], single-task zero-loss barrier merging [1] highlights the immense potential of unit matching in combining models from multiple task-specific tasks. Current unit matching methods primarily consider two types of matching spaces [27]: one is the activation space [1, 3, 28, 33, 40, 42–44], and the other is the weight space [1, 24, 35, 48, 51, 59]. Nevertheless, we demonstrate through a straightforward experiment that the similarity of units in the activation space and weight space is not always consistent, and these discrepancies manifest various pros and cons during model merging. Furthermore, simultaneously considering the similarity of models in both weight space and activation space can offer potential advantages in reducing merging barriers. Therefore, our method simultaneously constrains unit matching in both the weight space and activation space.

### 3. Methodology

In Sec. 3.1, we introduce a general procedure of merging model via unit matching. Afterward, we provide the theoretical basis for merging models under dual-space constraints and then elaborate on how to merge models under dual-space constraints in Sec 3.2. Two solutions for adaption to group structures are presented in Sec. 3.3.

#### 3.1. Preliminaries

Assume there are  $N$  pre-trained models in the same architecture. The  $n$ -th model is parameterized by weights  $\{\mathbf{W}_l^{(n)}\}_{l=1}^L$ , where  $L$  represents the number of layers of the models. Without loss of generality, we omit the bias term in parameters and simply regard the model as MLP model, so  $\mathbf{W}_l^{(n)} \in \mathbb{R}^{D_l D_{l-1}}$ , where  $D_l$  and  $D_{l-1}$  represent the number of units in  $l$ -th layer and  $(l-1)$ -th layer, respectively.

The main idea behind model merging revolves around the fusion of units that exhibit significant similarity, aiming to maximize the cumulative global similarity of the merged units within each layer. The unit parameters after merging are obtained by averaging the unit parameters before merging. The solution can be summarized as computing the maximum-weighted matching of the graph, where the nodes refer to the units and the edges refer to their similarities. The goal is to find a subset of edges with the maximum weight, and in which no node occurs more than  $N$  times. Formally, in the  $l$ -th layer, let  $\mathbb{C}_l \in \mathbb{R}^{(N \cdot D_l) \times (N \cdot D_l)}$  be the similarity matrix of merging a pair of units in the  $N$  models, where each element  $(\mathbb{C}_l)_{i,j}$  denotes the similarity of the  $i$ -th unit and the  $j$ -th unit, *i.e.*, the weights of edges as mentioned above. Then the goal of model merging can be formulated as follows:

$$\begin{aligned} \arg \max_{\mathbb{P}_l} \sum_{l=1}^L \langle \mathbb{C}_l, \mathbb{P}_l \rangle, \text{ s.t. } \sum_{i=1}^{N \cdot D_l} (\mathbb{P}_l)_{i,j} \leq N \quad \forall j, \\ (\mathbb{P}_l)_{i,j} = 1 \implies (\mathbb{P}_l)_{i,k} = (\mathbb{P}_l)_{j,k}, \quad \forall i, j, k \\ (\mathbb{P}_l)_{i,i} = 1 \quad \forall i. \end{aligned} \quad (1)$$

$\mathbb{P}_l$  denotes the merging matrix which is symmetrical and binary, and where  $(\mathbb{P}_l)_{i,j} = 1$  represents that the  $i$ -th unit and  $j$ -th unit are merged together.  $\langle \mathbf{X}, \mathbf{Y} \rangle = \sum_{i,j} \mathbf{X}_{i,j} \mathbf{Y}_{i,j}$  denotes the Frobenius inner product between real-valued matrices  $\mathbf{X}$  and  $\mathbf{Y}$ .

To merge the similar units from different models, we can construct a set of permutation matrices  $\{\mathbf{P}_l\}_{l=1}^L$ . Then the merge operation can be formulated as follows:

$$\mathbf{W}_l' = \frac{1}{N} \sum_{n=1}^N (\mathbf{P}_l^{(n)})^\top \mathbf{W}_l^{(n)} \mathbf{P}_l^{(n)}, \quad l \in \{1, 2, \dots, L\} \quad (2)$$

To clarify the relation between  $\mathbf{P}_l^{(n)}$  and  $\mathbb{P}_l$ , we define  $\mathbf{P}_l = \mathbf{P}_l^{(1)} \circ \mathbf{P}_l^{(2)} \circ \dots \circ \mathbf{P}_l^{(N)}$  where  $\mathbf{P}_l^{(n)} \in \mathbb{R}^{D_l D_l}$  and operation  $\circ$  denotes concatenating the matrices in the first dimension. Then  $\mathbf{P}_l$  can be straightforwardly derived from  $\mathbb{P}_l$  by removing duplicate columns of  $\mathbb{P}_l$  because of the constraints of Eq. (1).

#### 3.2. Merging under Dual-Space Constraints

As introduced in Section 1, prior works usually measure the unit similarity within a single space (*i.e.* the activation space or the weight space), resulting in inconsistencies between the two spaces. In this work, we assess unit similarity by considering similarities in both activation space and weight space. In this study, we propose an innovative merging framework that addresses disparities between models in both the activation space and weight space. The fundamental challenge lies in reconciling the distinctions between weight-based matching and activation-based matching during the matching procedure.

Given the  $\{\mathbf{A}_l\}_{l=1}^L$  and  $\{\mathbf{Z}_l\}_{l=1}^L$  as the representation vectors of the  $l$ -th layer in the activation space and the weight space, respectively, and with a balanced factor  $\alpha \in [0, 1]$ , the expected objective can be written as follows:

$$\arg \max_{\mathbb{P}_l} \alpha \sum_{l=1}^L \langle \mathbb{C}(\mathbf{Z}_l), \mathbb{P}_l \rangle + (1 - \alpha) \sum_{l=1}^L \langle \mathbb{C}(\mathbf{A}_l), \mathbb{P}_l \rangle \quad (3)$$

Here,  $\mathbb{C} : \mathbb{R}^{(N \cdot D_l) R_l} \rightarrow \mathbb{R}^{(N \cdot D_l)(N \cdot D_l)}$  denote the function that calculate the similarity matrix by  $\mathbf{A}_l$  or  $\mathbf{Z}_l$ . Specially,  $R_l$  refers to the quantity of related parameters when it comes to  $\mathbf{Z}_l$ , or  $R_l$  training data points when it comes to  $\mathbf{A}_l$ . We omit the constraint conditions here. Then we can re-express Eq. (3) as follows:

$$\arg \max_{\mathbb{P}_l} \sum_{l=1}^L \langle \alpha \mathbb{C}(\mathbf{Z}_l) + (1 - \alpha) \mathbb{C}(\mathbf{A}_l), \mathbb{P}_l \rangle \quad (4)$$

Manifestly, the aforementioned derivation demonstrates that the dual-space constraints can be realized through a linear combination of the similarity matrices in both spaces. This formulation presents an efficacious approach to address unit matching under the dual-space constraints. The remaining question is how to implement Eq. (4) in the matching procedure. To this end, we delve into the solutions within two distinct spaces.

In term of the activations, activation-based matching can be solved in a single-pass. This procedure is described as:

$$\mathbf{P}_l = \Psi(\mathbb{C}(\mathbf{A}_l)), l \in \{1, 2, \dots, L\} \quad (5)$$

Here,  $\Psi(\cdot)$  is the function that derives the permutation matrix by Eq. (1) from the similarity matrix.

In the realm of weight space, there are three distinct perspectives on how to obtain representation vectors. One viewpoint advocates the use of output vectors in the next layer to represent the units [59]. Another approach leans towards utilizing the weights of the current layer [41, 48]. An advanced solution involves considering all weights associated with the layer, including those from the preceding and subsequent layers. Such a method has achieved zero-loss barriers in the merging of single-task models [1]. In our work, we adopted the third method to obtain representations of units in the weight space. Consequently, for  $n$ -th model, the representation vectors in the  $l$ -th layer is conducted by:

$$\mathbf{Z}_l^{(n)} = \mathbf{W}_l^{(n)} \mathbf{P}_{l-1}^{(n)} \parallel \mathbf{P}_{l+1}^{(n)} (\mathbf{W}_{l+1}^{(n)})^\top \quad (6)$$

where  $\parallel$  denote the operation which concatenates the given matrices in the second dimension. Note that aforementioned  $\mathbf{W}_l$  and  $\mathbf{Z}_l$  contain the units of all models in the  $l$ -th layer like  $\mathbf{P}_l$  and can be expressed as  $\mathbf{W}_l^{(1)} \circ \dots \circ \mathbf{W}_l^{(N)}$  and  $\mathbf{Z}_l^{(1)} \circ \dots \circ \mathbf{Z}_l^{(N)}$ , respectively.

---

### Algorithm 1 MuDSC Merging

---

**Require:** The weights  $\{\mathbf{W}_l\}_{l=1}^L$  and the activations  $\{\mathbf{A}_l\}_{l=1}^L$  of the models, the function for computing the similarity matrix  $\mathbb{C}$ , selected matching algorithm  $\Psi$ , the balanced factor  $\alpha$ .  $RP$  returns a random permutation of the sequence.

**for**  $l = 1, 2, \dots, L$  **do**

$\mathbf{C}_l' \leftarrow \mathbb{C}(\mathbf{A}_l)$

$\mathbf{P}_l \leftarrow \Psi(\mathbf{C}_l')$

**end for**

**repeat**

**for**  $l = RP(1, 2, \dots, L)$  **do**

$\mathbf{Z}_l \leftarrow v(\mathbf{P}_{l-1}, \mathbf{P}_{l+1}, \mathbf{W}_l, \mathbf{W}_{l+1})$

$\mathbf{C}_l' \leftarrow \alpha \mathbb{C}(\mathbf{Z}_l) + (1 - \alpha) \mathbf{C}_l'$

$\mathbf{P}_l \leftarrow \Psi(\mathbf{C}_l')$

**end for**

**until** convergence

Get the merged weights  $\{\mathbf{W}_l'\}_{l=1}^L$  by Eq. (2)

**return**  $\{\mathbf{W}_l'\}_{l=1}^L$

---

There is a contradiction that a valid  $\mathbf{Z}_l$  for matching in  $l$ -th layer requires the units in the  $(l + 1)$ -th and  $(l - 1)$ -th layers to reach the optimal match. This contradiction can be resolved through an iterative algorithm, and its recursive expression is as follows:

$$\begin{cases} \mathbf{P}_l^1 = \mathbb{I} \\ \mathbf{Z}_l^{t+1} = v(\mathbf{P}_{l-1}^t, \mathbf{P}_{l+1}^t, \mathbf{W}_l^t, \mathbf{W}_{l+1}^t) \\ \mathbf{P}_l^{t+1} = \Psi(\mathbb{C}(\mathbf{Z}_l^{t+1})) \end{cases}, \quad (7)$$

where  $l \in \{1, 2, \dots, L\}$  and  $t \in \mathbb{N}^+$ . When the number of iterations approaches infinity, the model reaches its optimal match in weight space. In Eq. (7),  $\mathbb{I}$  is a permutation matrix where the submatrix of each model is identity matrix. And  $v(\cdot, \cdot, \cdot, \cdot)$  is the function that applies permutation matrices to the weights of all models by Eq. (6) and then concatenate all the  $\mathbf{Z}_l^{(n)}$  into  $\mathbf{Z}_l$ .

Afterwards, it is evident to implement the merging method under the dual-space constraints based on Eq. (4). We substitute Eq. (4) and Eq. (5) into Eq. (7), yielding:

$$\begin{cases} \mathbf{P}_l^1 = \Psi(\mathbb{C}(\mathbf{A}_l)) \\ \mathbf{Z}_l^{t+1} = v(\mathbf{P}_{l-1}^t, \mathbf{P}_{l+1}^t, \mathbf{W}_l^t, \mathbf{W}_{l+1}^t) \\ \mathbf{P}_l^{t+1} = \Psi(\alpha \mathbb{C}(\mathbf{Z}_l^{t+1}) + (1 - \alpha) \mathbb{C}(\mathbf{A}_l)) \end{cases}, \quad (8)$$

Lastly, we present the overall procedure of our MuDSC merging in Alg. 1. We further illustrate the convergence and the complexity of Alg. 1 in the supplementary material.

### 3.3. Adaptation to Group Structure

As mentioned previously, the abstractions of current matching algorithm do not generalize to networks with the following structure:

1. **External Permutation Invariance:** There are  $g$  separate groups which can be permuted with each other. For instance, Group Normalization [50] divides channels into  $g$  groups and normalizes the features within each group.
2. **Internal Permutation Invariance:** Each of those  $g$  groups contains  $k$  units which can be permuted. For instance, the features within each group of Group Normalization can be permuted.

This structure is commonly encountered within contemporary neural network modules, in particularly Group Normalization [50] and Multi-Head Attention [47]. Consequently the application of matching algorithms will face significant hindrance if they cannot adapt to this structure. Here, we propose viable adaptation strategies for an alignment algorithm based on linear sum assignment [1] (namely group alignment) and the greedy algorithm used in zip operation [43] (namely group zip).

Among them, group alignment first performs internal alignment of pairwise groups and then conducts external alignment based on the average similarity obtained from the internal alignment. We provide an illustrative diagram in Figure 2. While this algorithm has quadratic time complexity when it comes to external matching, we have found this cost to be acceptable for the application in our experiments. In addition, it is worth noting that this algorithm is capable of achieving a global optimal match.




Figure 2. **Left.** Examples of vanilla structure and group structure. **Right.** An example of group alignment. First, we calculate the similarity between units. Next, we compute permutation and then calculate the average of matched similarity within each pairs of groups. Finally, we compute permutation for each pairs of groups and then set the permutation of unmatched pairs to zeros.

For group zip, unlike alignment-based approaches, zip operation employs a greedy matching algorithm to merge similar units. However, this method is challenging to extend to group merging. In response to this, we propose a simplified matching scheme. We initially determine which groups to merge together through a group matching strategy, and then we "zip" these groups individually. As finding a

suitable grouping strategy is a challenging task, we did not introduce an efficient grouping strategy in this work, leaving the task for future work. In subsequent experiments, we directly use the group matching results of the group alignment as the outcome of our grouping strategy, enabling a fair performance comparison between the alignment methods and the zip methods.

## 4. Experiments

We first validate the superiority of MuDSC to existing methods in merging models of homogeneous tasks (*i.e.*, classification tasks with different sets of categories) on both the small- and the large-scale datasets (Section 4.1). Then, we verify the effectiveness of MuDSC in merging models of heterogeneous tasks (*e.g.*, segmentation and classification) with the pre-trained models released in Taskonomy [60] (Section 4.2). Finally, we visualize multi-task loss landscape to further demonstrate the superiority of MuDSC over prior single-space methods (Section 4.3).

### 4.1. Merging Models of Homogeneous Tasks

#### 4.1.1 Experimental Settings

**Datasets and models** The experiments are conducted on the small-scale CIFAR-10 [23], CIFAR-100 [23] and the large-scale ImageNet [8]. In the selection of models for merging, we emphasize the initial state of the model and whether it has a group structure *i.e.* Multi-Head Attention [9] and Group Normalization [50]. These two aspects of configuration significantly reflect the effectiveness and applicability of the merging algorithm. For randomly initialized models, ResNet-20 [16] and ResNet-20 with Group Normalization [50] are selected. Additionally, we choose several pre-trained models from PyTorch Image Models (timm) [49], including ResNet-26 [16], ResNet-50 with Group Normalization [50], ViT-Small [9], DINO-Small [5] and Swin-Tiny [29].

**Baselines.** Our method is applied to two different types of matching algorithms referred to as alignment-based matching and zip-based matching. We compare our method with the current state-of-the-art methods of these two types. For alignment-based matching, the algorithm aligns the units of all models with those of the first model among them [1]. We implement an alignment-based version of our method  $\text{MuDSC}_{Align}$  and compare it with Git Rebasin [1]. Additionally, we incorporated an activation-based alignment algorithm (abbreviated as A. Align). For zip-based matching, the units from different models are "zipped" together in the same layer which is referred to as zip operation [43]. Similarly, we implement  $\text{MuDSC}_{Zip}$  based on zip operation, comparing it to Zipit [43] and a weight-based variant (abbreviated as W.Zip) implemented by us.

**Experimental details.** We randomly partition a classification dataset into two non-overlapping sub-classification

Model	Resnet20								Resnet20GN			
Dataset	CIFAR100(50+50)				CIFAR10(5+5)				CIFAR100(50+50)			
Method	Joint	Avg	T. A	T. B	Joint	Avg	T. A	T. B	Joint	Avg	T. A	T. B
Average	16.52	24.22	23.22	25.21	54.42	75.24	79.58	70.90	5.63	11.06	9.67	12.44
Rebasin	41.33	56.94	57.31	56.58	60.61	88.57	88.46	88.68	13.85	22.18	22.99	21.37
A. Align	44.33	61.13	61.61	60.66	<b>61.71</b>	89.21	88.63	<b>89.78</b>	29.37	42.05	41.05	43.05
MuDSC <sub>Align</sub>	<b>45.50</b>	<b>62.81</b>	<b>63.06</b>	<b>62.56</b>	60.84	<b>89.34</b>	<b>89.04</b>	89.63	<b>31.84</b>	<b>45.31</b>	<b>45.34</b>	<b>45.29</b>
Zipit	54.69	66.78	67.11	66.44	82.44	94.61	94.22	95.00	29.93	41.20	39.99	42.41
W.Zip	55.16	67.65	68.58	66.71	82.85	94.71	94.42	94.99	14.28	20.95	19.17	22.72
MuDSC <sub>Zip</sub>	<b>56.01</b>	<b>68.13</b>	<b>68.80</b>	<b>67.47</b>	<b>83.09</b>	<b>94.88</b>	<b>94.56</b>	<b>95.21</b>	<b>30.05</b>	<b>41.52</b>	<b>40.39</b>	<b>42.65</b>

Table 1. The joint accuracy and the per-task accuracy of the merged multitask model. Two original models are trained from scratch on two subtasks of the same classification task. We emphasize the data achieving the best accuracy. Std is provided in the supplementary material.

Model	Resnet26				Resnet50GN				ViT			
Method	Joint	Avg	T. A	T. B	Joint	Avg	T. A	T. B	Joint	Avg	T. A	T. B
Average	61.44	74.75	74.46	75.05	74.52	84.78	85.06	84.50	70.16	84.32	84.32	84.32
Rebasin	61.39	74.79	74.48	75.10	74.52	84.78	85.06	84.50	<b>70.16</b>	84.32	84.32	84.32
A. Align	61.91	75.41	75.03	75.79	74.44	84.77	84.99	84.56	69.99	84.22	84.20	84.24
MuDSC <sub>Align</sub>	<b>62.84</b>	<b>76.14</b>	<b>75.87</b>	<b>76.40</b>	<b>74.66</b>	<b>84.91</b>	<b>85.25</b>	<b>84.58</b>	70.09	<b>84.39</b>	<b>84.38</b>	<b>84.40</b>
Zipit	60.23	73.68	73.20	74.17	72.05	82.99	83.06	82.92	68.57	83.05	82.79	83.30
W.Zip	61.28	74.69	74.42	74.96	74.52	84.78	85.06	84.50	<b>70.16</b>	84.32	84.32	84.32
MuDSC <sub>Zip</sub>	<b>61.58</b>	<b>75.01</b>	<b>74.61</b>	<b>75.41</b>	<b>74.71</b>	<b>84.88</b>	<b>85.14</b>	<b>84.62</b>	70.10	<b>84.38</b>	<b>84.41</b>	<b>84.36</b>

Table 2. The joint accuracy and the per-task accuracy of the merged multitask model. Two original models that well-pretrained on ImageNet are finetuned on two subtasks of CIFAR100. We emphasize the data achieving the best accuracy. Std is in the supplementary material.

Method	Joint Acc	Avg Acc	T. A	T. B
Average	44.85	62.31	61.93	62.68
Rebasin	44.85	62.31	61.93	62.68
A. Align	44.86	62.62	62.56	62.67
MuDSC <sub>Align</sub>	<b>44.87</b>	<b>62.66</b>	<b>62.57</b>	<b>62.74</b>
Zipit	44.22	62.25	62.23	62.26
W.Zip	44.85	62.31	61.93	62.68
MuDSC <sub>Zip</sub>	<b>44.86</b>	<b>62.59</b>	<b>62.45</b>	<b>62.72</b>

Table 3. Results on ImageNet400(200+200). Merging ResNet-50 models trained with DINO pretrained backbone on disjoint 200 category subsets (Task A and B) of ImageNet-1k.

tasks, trained respective models for each, and subsequently merged the models into one. Then we evaluate performance of merged model with joint accuracy and per-task accuracy. Joint accuracy is the overall accuracy of a model when it is evaluated on all classes within a combined dataset. It is similar to a continual learning setting where our objective is to enhance the knowledge of the model. For per-task accuracy, we provided the accuracy of the merged multi-task model on two individual tasks, along with their average per-

formance. Each model is trained with a CLIP-style loss [36] using CLIP text encodings of the class names as targets. For fair comparisons, we train 5 pairs of models and report the average accuracy. Three zip-based methods share the best hyperparameters (mentioned in [43]) on Zipit. In addition, for each pair of models, we search the balance factor  $\alpha$  of MuDSC for the best per-task accuracy on the training data and then apply it to test.

#### 4.1.2 Results and Analysis

**Results on CIFAR.** We first test the performance of the methods on merging randomly initialized models in Table 1. For experiments with Resnet20 on CIFAR100 and CIFAR10, the zip-based methods achieve relatively higher merged accuracy than alignment-based methods and our MuDSC<sub>Zip</sub> achieve further improvements. In experiments with Resnet20GN, MuDSC<sub>Align</sub> outperform all methods and significantly surpass the second-best method by 8.41%, 7.75%, 10.45%, 5.20% on joint accuracy and three per-task accuracy. Secondly, we conduct the experiments on CIFAR100 with Resnet26, Resnet50GN and ViT all of which are well-pretrained on ImageNet [8]. Unlike randomly initialized models, finetuned models exhibit a high similarity in weight space as they are initialized with the same pretrained

Visual Task	Method						
	Weight Average	Rebasin	Act. Align	MuDSC <sub>Align</sub>	Zipit	Weight Zip	MuDSC <sub>Zip</sub>
Class Object	76.15	80.76	<b>89.75</b>	<b>89.75</b> +0.00	84.84	80.93	<b>86.04</b> +1.20
Segment Semantic	23.01	30.59	52.54	<b>55.24</b> +2.69	32.51	33.77	<b>36.22</b> +2.45
Rbg2depth	95.42	95.90	98.69	<b>98.70</b> +0.01	99.30	99.07	<b>99.35</b> +0.05
Rbg2mist	94.56	95.00	<b>98.28</b>	<b>98.28</b> +0.00	99.08	98.59	<b>99.15</b> +0.00
Edge3D	79.79	80.44	<b>91.04</b>	90.77 -0.27	<b>93.22</b>	86.42	92.81 -0.40
Edge2D	68.54	75.48	<b>81.37</b>	81.24 -0.12	90.39	88.85	<b>93.09</b> +2.69
Keypoints2D	71.33	75.89	81.99	<b>82.30</b> +0.31	93.38	93.62	<b>94.73</b> +1.11
Keypoints3D	91.54	91.99	96.16	<b>96.17</b> +0.02	<b>96.99</b>	96.00	96.95 -0.04
Reshading	-4.72	8.52	60.23	<b>61.85</b> +1.61	<b>69.18</b>	41.66	68.71 -0.47
Rbg2sfnorm	26.62	29.43	65.46	<b>65.56</b> +0.09	76.62	64.43	<b>77.50</b> +0.88
Autoencoding	21.96	35.65	51.84	<b>54.21</b> +2.37	86.55	77.58	<b>87.99</b> +1.45
Denoising	33.92	33.57	52.37	<b>54.40</b> +2.03	84.28	76.14	<b>85.49</b> +1.22
<b>Total Average</b>	56.51	61.10	76.64	<b>77.37</b> +0.73	83.86	78.09	<b>84.84</b> +0.98

Table 4. The average scaled performance of multiple heterogenous-task models after merged with each other. We emphasize the data achieving the best accuracy and the improvement of our methods.

model, making it challenging to discover better matching results. As observed in Table 2, the results for W. Zip and Rebasin are completely consistent with those of Average. Another regrettable fact is that when merging finetuned models, activation-based methods perform even worse than direct averaging. Fortunately, from Table 2, we can see that merging models under the dual-space constraints has overcome the obstacles of the aforementioned single-space methods, achieving further improvements in merged accuracy. Results of DINO-S and Swin-T (Table S7) also validate the superiority of our method.

**Results on ImageNet.** To test our method on the large scale data, we train 5 ResNet-50 models initialized with DINO pretrained backbone [5] on disjoint 200 class subsets of ImageNet-1k [8]. Then we conduct experiments on exhaustively merging pairs from the 5 models. As shown in Tab. 3, the proposed MuDSC still surpasses existing methods, demonstrating its capability for model merging on the large-scale data.

## 4.2. Merging Models of Heterogenous Tasks

### 4.2.1 Experimental Settings

**Tasks and models.** We adopt 12 pre-trained models from Taskonomy [60] trained on various tasks (including Autoencoder, Denoise, Edge 2D, Edge 3D, Keypoint 2D, Keypoint 3D, Reshade, Rbg2depth, Rbg2mist, Rbg2sfnorm, Segmentation, and Classification.). The architectures of these models follows an encoder-decoder scheme, in which the encoder is implemented by fully convolutional layers and the decoder varies according to the tasks. Please refer to [60] for more detailed information. In our experiments, only the encoders of these models are adopted for merging.

**Baselines and Metric.** The settings of baselines are the

same as that of Sec. 4.1. In order to reflect the gap between the merged model and the original model more intuitively, we define a scaled performance of model  $\theta$ , which is expressed as follows:

$$\mathcal{L}_{SP} = \frac{\mathcal{L}_\theta - \mathcal{L}_0}{\mathcal{L}_1 - \mathcal{L}_0}, \quad (9)$$

where  $\mathcal{L}_\theta$ ,  $\mathcal{L}_1$  and  $\mathcal{L}_0$  represent the original performance metrics for the model  $\theta$ , the pretrained model and the average estimator. Notably, when  $\mathcal{L}_{SP}$  approaches 1, the performance of the model  $\theta$  is similar to the pre-trained model, and when  $\mathcal{L}_{SP}$  approaches 0, the performance of the model  $\theta$  is similar to the average value of the labels. In other words, we aim for the performance of the merged model to be closer to, or even better than 1. For the original performance, we maintain consistency with [60] except for adopting error rate in Segmentation, and we are unable to access any GAN loss due to the lack of discriminators.

**Experimental Details.** We merge the adopted pretrained models with each other and ultimately report the average scaled performance per and all tasks for all merging methods. Due to the remarkably large size of the complete Taskonomy dataset, we conduct a reasonable sampling. For evaluation, we used the test split of official partitioned tiny subset, which contains 54513 samples, to evaluate the performance of the models. In order to reset the batch normalization of models and obtain the activations of the model, we sampled data from training dataset, which included 8 buildings and a total of 61520 samples. Additionally, the entire encoder is used to match but only the layers before the last layer are merged. In this experiment, we simply adopted 0.5 as the balanced factor for our MuDSC.




Figure 3. The loss landscape visualization of three zip-based methods (MuDSC<sub>Zip</sub>, Zipit and W. Zip). Stars, squares, and diamond mark the positions of the merged models, and black spots mark the positions of their parent models.

#### 4.2.2 Results and Analysis

Tab. 4 provide a quantitative comparison between the proposed MuDSC with existing works in scaled performance. It can be easily seen that the MuDSC yields significantly superior performance to existing two types of matching methods. In alignment-based algorithms, MuDSC<sub>Align</sub> maintains or surpasses the current SOTA in 83% of tasks, Furthermore, in certain tasks *i.e.* Segment Semantic, Reshading, Autoencoding, and Denoise, MuDSC<sub>Align</sub> brings improvements of 5.12%, 2.68%, 4.58%, and 3.88%, respectively. Similarly, in zip-based methods, MuDSC outperforms the current SOTA methods in 75% of tasks, with improvements exceeding 1% in 8 tasks, notably achieving a remarkable 11.42% improvement in the Segment Semantic task. Overall, regardless of alignment-based methods and zip-based methods, MuDSC brings significant improvements to both.

#### 4.3. Loss Landscape Visualization

In this subsection, we demonstrate an example of how MuDSC improves the zip-based method. We conduct visualizations on a multi-task experiment constructed with MNIST [25]. Two binary classification tasks involve determining whether a sample is a prime number (*i.e.* Task A) and whether a sample is odd (*i.e.* Task B). The architecture of the model is a four-layer MLP and each model is trained as CLIP image encoder. The results are reported at Tab. S1. Then, we visualize merged models and their original models obtained by various methods onto a single diagram to intuitively compare their differences in the flatness of the loss landscape. To be more specific, we construct a set of high-dimensional vectors obtained by flattening models' parameters and use the PCA dimension reduction algorithm [32] to generate the two-dimensional coordinates. In Fig. 3, there are the average loss landscape, the loss landscape for Task A, and the loss landscape for Task B from left to right.

Then we mark the positions of activation-based methods (*i.e.* Zipit), weight-based methods, and MuDSC<sub>Zip</sub>.

As shown in Fig. 3, in the average loss landscape, our approach positions the merged model directly at the lowest point of the basin, in contrast to the activation-based method which places the merged model near the lowest point, and the weight-based method, which positions the merged model further away from the center of the basin. Then observe the loss landscape of different tasks. Upon further examination of the loss landscapes for different tasks, we observe that MuDSC facilitates the matching algorithm to discover a set of overlapping loss basins at lower points. This implies that our approach can better balance the performance of different tasks when merging for multi-task scenarios. Overall, with the support of a more precise constraint, MuDSC achieves an enhancement of the matching algorithm.

### 5. Conclusion

We introduce Merging under Dual-Space Constraints (MuDSC) to balance the inconsistency of unit similarity in weight space and activation space when merging models. MuDSC linearly combines the similarity matrices both of the weights and the activations of the units to seek a better permutation matrix. We find experimentally that MuDSC enhances the performance of the merged multi-task model across various tasks and architecture. We conduct the visualization of the merged model in the multi-task loss landscape which shows that MuDSC makes the merged model locate in the overlapping part with unified lower loss of each task.

### 6. Acknowledgements

This work is supported by the Science and Technology Project of SGCC: Research and Digital Application of High-precision Electric Power Super-scale Pre-trained Visual Model (5108-202218280A-2-395-XG).



## References

- [1] Samuel Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models modulo permutation symmetries. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 3, 4, 5
- [2] Yamini Bansal, Preetum Nakkiran, and Boaz Barak. Re-visiting model stitching to compare neural representations. *Advances in neural information processing systems*, 34:225–236, 2021. 2
- [3] Frederik Benzing, Simon Schug, Robert Meier, Johannes Von Oswald, Yassir Akram, Nicolas Zucchet, Laurence Aitchison, and Angelika Steger. Random initialisations performing above chance and how to find them. *arXiv preprint arXiv:2209.07509*, 2022. 3
- [4] Johann Brea, Berfin Simsek, Bernd Illing, and Wulfram Gerstner. Weight-space symmetry in deep networks gives rise to permutation saddles, connected by equal-loss valleys across the loss landscape. *arXiv preprint arXiv:1907.02911*, 2019. 3
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 5, 7
- [6] Adrián Csiszárík, Péter Kőrösi-Szabó, Akos Matszangosz, Gergely Papp, and Dániel Varga. Similarity and matching of neural network representations. *Advances in Neural Information Processing Systems*, 34:5656–5668, 2021. 2
- [7] Marcus de Carvalho, Mahardhika Pratama, Jie Zhang, and Yajuan Sun. Class-incremental learning via knowledge amalgamation. In *Machine Learning and Knowledge Discovery in Databases*, pages 36–50, Cham, 2023. Springer Nature Switzerland. 1
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 5, 6, 7
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 2, 5
- [10] Rahim Entezari, Hanie Sedghi, Olga Saukh, and Behnam Neyshabur. The role of permutation invariance in linear mode connectivity of neural networks. *arXiv preprint arXiv:2110.06296*, 2021. 2, 3
- [11] Chris Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, and Chelsea Finn. Efficiently identifying task groupings for multi-task learning. *Advances in Neural Information Processing Systems*, 34:27503–27516, 2021. 1
- [12] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018. 2
- [13] K. Fukumizu and S. Amari. Local minima and plateaus in hierarchical structures of multilayer perceptrons. *Neural Networks*, 13(3):317–327, 2000. 3
- [14] Shangde Gao, Yichao Fu, Ke Liu, and Yuqiang Han. Contrastive knowledge amalgamation for unsupervised image classification. In *Artificial Neural Networks and Machine Learning – ICANN 2023*, pages 192–204, Cham, 2023. Springer Nature Switzerland. 1, 2
- [15] Elisenda Grigsby, Kathryn Lindsey, and David Rolnick. Hidden symmetries of ReLU networks. In *Proceedings of the 40th International Conference on Machine Learning*, pages 11734–11760. PMLR, 2023. 3
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 5
- [17] Robert Hecht-Nielsen. On the algebraic structure of feedforward network weight spaces. In *Advanced Neural Computers*, pages 129–135. North-Holland, Amsterdam, 1990. 3
- [18] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 3
- [19] Xisen Jin, Xiang Ren, Daniel Preotiuc-Pietro, and Pengxiang Cheng. Dataless knowledge fusion by merging weights of language models. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [20] Yongcheng Jing, Yiding Yang, Xinchao Wang, Mingli Song, and Dacheng Tao. Amalgamating knowledge from heterogeneous graph neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [21] Yongcheng Jing, Chongbin Yuan, Li Ju, Yiding Yang, Xinchao Wang, and Dacheng Tao. Deep graph reprogramming. In *CVPR*, 2023. 2
- [22] Keller Jordan, Hanie Sedghi, Olga Saukh, Rahim Entezari, and Behnam Neyshabur. REPAIR: RENormalizing permuted activations for interpolation repair. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 3
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [24] Thanh Chi Lam, Nghia Hoang, Bryan Kian Hsiang Low, and Patrick Jaillet. Model fusion for personalized learning. In *Proceedings of the 38th International Conference on Machine Learning*, pages 5948–5958. PMLR, 2021. 3
- [25] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010. 8
- [26] Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 991–999, 2015. 2
- [27] Weishi Li, Yong Peng, Miao Zhang, Liang Ding, Han Hu, and Li Shen. Deep model fusion: A survey. *arXiv preprint arXiv:2309.15698*, 2023. 2, 3
- [28] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. Convergent learning: Do different neural

- networks learn the same representations? *arXiv preprint arXiv:1511.07543*, 2015. 2, 3
- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 5
- [30] Sihui Luo, Xinchao Wang, Gongfan Fang, Yao Hu, Dapeng Tao, and Mingli Song. Knowledge amalgamation from heterogeneous networks by common feature learning. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. 1, 2
- [31] Sihui Luo, Wenwen Pan, Xinchao Wang, Dazhou Wang, Haihong Tang, and Mingli Song. Collaboration by competition: Self-coordinated knowledge amalgamation for multi-talent student learning. In *European Conference on Computer Vision*, 2020. 1, 2
- [32] Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (pca). *Computers & Geosciences*, 19(3): 303–342, 1993. 8
- [33] Dang Nguyen, Trang Nguyen, Khai Nguyen, Dinh Phung, Hung Bui, and Nhat Ho. On cross-layer alignment for model fusion of heterogeneous neural networks. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023. 3
- [34] Zizheng Pan, Jianfei Cai, and Bohan Zhuang. Stitchable neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16102–16112, 2023. 2
- [35] Fidel A Guerrero Peña, Heitor Rapela Medeiros, Thomas Dubail, Masih Aminbeidokhti, Eric Granger, and Marco Pedersoli. Re-basin via implicit sinkhorn differentiation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20237–20246, 2023. 3
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6
- [37] Chengchao Shen, Xinchao Wang, Jie Song, Li Sun, and Mingli Song. Amalgamating knowledge towards comprehensive classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3068–3075, 2019. 1, 2
- [38] Chengchao Shen, Mengqi Xue, Xinchao Wang, Jie Song, Li Sun, and Mingli Song. Customizing student networks from heterogeneous teachers via adaptive knowledge amalgamation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3504–3513, 2019. 2
- [39] Berfin Simsek, François Ged, Arthur Jacot, Francesco Spadaro, Clément Hongler, Wulfram Gerstner, and Johanni Brea. Geometry of the loss landscape in overparameterized neural networks: Symmetries and invariances. In *International Conference on Machine Learning*, pages 9722–9732. PMLR, 2021. 3
- [40] Sidak Pal Singh and Martin Jaggi. Model fusion via optimal transport. *Advances in Neural Information Processing Systems*, 33, 2020. 2, 3
- [41] Sidak Pal Singh and Martin Jaggi. Model fusion via optimal transport. *Advances in Neural Information Processing Systems*, 33, 2020. 4
- [42] Jie Song, Zhengqi Xu, Sai Wu, Gang Chen, and Mingli Song. Modelgif: Gradient fields for model functional distance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 3
- [43] George Stoica, Daniel Bolya, Jakob Bjorner, Taylor Hearn, and Judy Hoffman. Zipit! merging models from different tasks without training, 2023. 2, 5, 6
- [44] Norman Tatro, Pin-Yu Chen, Payel Das, Igor Melnyk, Prasanna Sattigeri, and Rongjie Lai. Optimizing mode connectivity via neuron alignment. *Advances in Neural Information Processing Systems*, 33:15300–15311, 2020. 2, 3
- [45] Jidapa Thadajarassiri, Thomas Hartvigsen, Xiangnan Kong, and Elke A Rundensteiner. Semi-supervised knowledge amalgamation for sequence classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(11):9859–9867, 2021. 1
- [46] Jidapa Thadajarassiri, Thomas Hartvigsen, Walter Gerych, Xiangnan Kong, and Elke Rundensteiner. Knowledge amalgamation for multi-label classification via label dependency transfer. 37:9980–9988, 2023. 2
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [48] Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni. Federated learning with matched averaging. In *International Conference on Learning Representations*, 2020. 3, 4
- [49] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. 5
- [50] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2, 5
- [51] Peng Xiao, Biao Zhang, Samuel Cheng, Ke Wei, and Shuqin Zhang. Probabilistic fusion of neural networks that incorporates global information. In *Proceedings of The 14th Asian Conference on Machine Learning*, pages 1149–1164. PMLR, 2023. 3
- [52] Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. Resolving interference when merging models. In *NeurIPS*, New Orleans, USA, 2023. Proceedings of Machine Learning Research. 2, 3
- [53] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. *arXiv preprint arXiv:2310.02575*, 2023. 2, 3
- [54] Xingyi Yang, Jingwen Ye, and Xinchao Wang. Factorizing knowledge in neural networks. In *European Conference on Computer Vision*, pages 73–91. Springer, 2022. 2
- [55] Xingyi Yang, Daquan Zhou, Songhua Liu, Jingwen Ye, and Xinchao Wang. Deep model reassembly. *Advances in neural information processing systems*, 35:25739–25753, 2022. 2

- [56] Jingwen Ye, Yixin Ji, Xinchao Wang, Kairi Ou, Dapeng Tao, and Mingli Song. Student becoming the master: Knowledge amalgamation for joint scene parsing, depth estimation, and more. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2829–2838, 2019. [2](#)
- [57] Jingwen Ye, Xinchao Wang, Yixin Ji, Kairi Ou, and Mingli Song. Amalgamating filtered knowledge: Learning task-customized student from multi-task teachers. In *International Joint Conference on Artificial Intelligence*, 2019. [2](#)
- [58] Jingwen Ye, Yixin Ji, Xinchao Wang, Xin Gao, and Mingli Song. Data-free knowledge amalgamation via group-stack dual-gan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12516–12525, 2020. [2](#)
- [59] Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 7252–7261. PMLR, 2019. [3](#), [4](#)
- [60] Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [5](#), [7](#)
- [61] Haofei Zhang, Feng Mao, Mengqi Xue, Gongfan Fang, Zunlei Feng, Jie Song, and Mingli Song. Knowledge amalgamation for object detection with transformers. *IEEE Transactions on Image Processing*, 32:2093–2106, 2023. [2](#)
- [62] Jinghan Zhang, Shiqi Chen, Junteng Liu, and Junxian He. Composing parameter-efficient modules with arithmetic operations. *arXiv preprint arXiv:2306.14870*, 2023. [3](#)
- [63] Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609, 2021. [1](#)