

Bridging Domains through Subspace-Aware Model Merging

Levy Chaves¹, Chao Zhou², Rebekka Burkholz², Eduardo Valle³, Sandra Avila¹

¹ Universidade Estadual de Campinas (UNICAMP), Recod.ai Lab., Instituto de Computação, Brasil

² CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

³ Intercom

Abstract

Model merging integrates multiple task-specific models into a single consolidated one. Recent research has made progress in improving merging performance for in-distribution or multi-task scenarios, but domain generalization in model merging remains underexplored. We investigate how merging models fine-tuned on distinct domains affects generalization to unseen domains. Through an analysis of parameter competition in the task matrix using singular value decomposition, we show that merging models trained under different distribution shifts induces stronger conflicts between their subspaces compared to traditional multi-task settings. To mitigate this issue, we propose SCORE (Subspace CONflict-Resolving mERging), a method designed to alleviate such singular subspace conflicts. SCORE finds a shared orthogonal basis by computing the principal components of the concatenated leading singular vectors of all models. It then projects each task matrix into the shared basis, pruning off-diagonal components to remove conflicting singular directions. SCORE consistently outperforms, on average, existing model merging approaches in domain generalization settings across a variety of architectures and model scales, demonstrating its effectiveness and scalability.

1. Introduction

The widespread availability of pre-trained models and public repositories has driven the development of techniques for efficiently combining and reusing existing models. Among these techniques, model merging [47] combines multiple fine-tuned models into a single unified model without requiring shared data or further fine-tuning. It directly merges the parameters of models fine-tuned from the same pre-trained backbone, enabling inference with a single forward pass, reducing both latency and storage. However, naive parameter averaging often performs poorly, as fine-tuning across tasks leads to divergent parameter distributions, resulting in interference and conflicts between task-specific

representations. Addressing these inconsistencies is crucial for effective and reliable model merging.

A popular strategy uses *task vectors* [14], defined as the difference between fine-tuned and pre-trained weights. TIES [45] identified redundancy in these vectors and proposed trimming small-magnitude elements while aggregating only those with consistent signs to mitigate conflicts. DARE [49] introduced stochastic sparsification, showing that up to 99% of model’s parameters can be removed in large models without notable degradation. More recently, Task Singular Vectors (TSV) [8] applied singular value decomposition (SVD) to task matrices at the layer level to measure task interference and merge parameters within the singular subspace.

Despite significant progress in model merging, most existing works primarily evaluate in-distribution performance, while the impact of model merging on domain generalization remains largely unexplored. Concretely, we treat each fine-tuned model as an expert on a particular distribution shift (a domain) and compose those experts to generalize to unseen domains. For intuition, consider two models fine-tuned on animals: one in rocky terrains, another in snowy land. Each model performs well in its domain, but neither captures both “rocky” and “wintry” attributes. By merging these models, we expect the merged model to combine complementary domain-specific representations and therefore to perform better on an unseen target such as animals in snowy mountain terrains, where these attributes coexist.

Analyzing the singular subspaces of such models, we observe subspace conflicts between fine-tunings on different domains that are much stronger than those of standard multi-task settings. We argue that such conflicts is challenging for SVD-based model merging and, to mitigate this issue, we propose SCORE, a Subspace CONflict-Resolving mERging method.

Our main contributions are:

1. We study the domain generalization of merged models using a leave-one-domain-out protocol.
2. We show that the overlap of task-matrix singular subspaces under domain generalization is significantly

higher than under multi-task scenarios.

3. We propose SCORE, a Subspace CONflict-Resolving mErging method that improves domain generalization.
4. Our extensive experiments show that SCORE consistently outperforms, on average, existing model merging methods across eight domain generalization benchmarks and three model scales.

2. Related Work

Model merging offers an efficient opportunity for combining multiple models without retraining. Several approaches rely on neuron permutations to align the models into a shared optimization basin through weight permutation strategies [1, 37, 44], or activations [16]. Another avenue focuses on the multi-task scenario, where a single pre-trained model is individually fine-tuned for different tasks and then combined for a multi-task model. Ilharco et al. [14] introduced the concept of *task vectors*, which are the parameter-wise difference between fine-tuned and pre-trained models, and the goal is to merge all available task vectors to create a multi-task model. TIES [45] addressed parameter redundancies, *i.e.*, when parameters benefit one task but not another, by first selecting the top- k most significant parameter changes and then constructing a sign vector based on the majority sign across all models. The latter is used to merge the task vectors disjointly, meaning the average is not computed when a parameter is zero or when parameters disagree in sign. Other alternatives to account for parameter redundancy include parameter sign consensus [42], randomly dropping and rescaling remaining parameters [49], keeping only the middle magnitudes [5], keeping only the highest magnitudes [26], or inter and cross-parameter competition to adjust the scale of individual parameters [10].

Another line of work mitigates parameter interference using Singular Value Decomposition (SVD) of individual weight matrices. KnOTs [38] concatenated task-specific low-rank adapters and averages right-singular vectors before SVD reconstruction, while Twin Merging [25] employs a router network to select among SVD experts. TSV [8] combined multiple SVD decompositions and orthogonalized the resulting singular vectors to reduce interference. ISO-C [27] enforced a uniform (isotropic) spectrum, improving multi-task performance.

Other methods rely on gradient-based optimization or validation-set information. Fisher Merging [28] and Reg-Mean [15] performed weighted averaging using the Fisher information matrix and input vector inner products, respectively. Ortiz-Jimenez et al. [31] demonstrated the benefits of NTK linearization for merging, while Yang et al. [46] learns layer-wise scalars through test-time training. EMR-Merging [13] further introduces per-task parameter masks and rescalers, assuming access to target fine-tuned weights.

Model merging for generalization. Wortsman et al. [43] showed that averaging parameters of multiple ImageNet fine-tuned models with different hyperparameter configurations can improve robustness against distribution shifts. Ramé et al. [35] improved previous work by requiring access to auxiliary tasks related to the target one, and relying on a two-step fine-tuning to recycle public models available for similar data as auxiliary tasks. While these methods show improvements, they rely on averaging strategies across hyper-parameters configurations or require access to validation datasets for hyperparameter selection.

For strategies beyond model averaging, Yang et al. [46] also considered corrupted versions of ImageNet as out-of-distribution evaluation, but their method relies on accessing test data samples to perform costly test-time training. However, as our scope is to evaluate the merged model without accessing any data or any optimization, our method only requires access to the fine-tuned source models and does not require any further training or gradient steps.

Guodong et al. [10] adopted a fixed set of generalization tasks for emotion classification in NLP tasks, without any further variation of such out-of-distribution datasets. Practitioners expect model-merging methods to be robust to several distribution shifts, and using a fixed set of datasets is suboptimal because we are only measuring performance for that particular case, making it unclear whether these methods generalize to other out-of-distribution data.

Current evaluation frameworks either compromise the advantages of model merging through data access requirements [35, 39, 46] or provide insufficient evidence of generalization across a fixed set of generalization tasks [10]. In contrast, our work addresses these concerns by requiring only access to fine-tuned source models and by considering merging models without any additional training, optimization, or privileged information. To the best of our knowledge, we are the first to investigate model merging under the domain generalization scheme of leave-one-out evaluation, providing a comprehensive evaluation across multiple distribution-shift scenarios.

3. Background and Motivation

In this section, we first outline the general framework for model merging and domain generalization, explaining how we connect these two concepts. Next, we introduce the notation used throughout the paper. Finally, we motivate our approach by measuring the subspace overlap between fine-tuned models across different domains, demonstrating that this overlap is greater than what is typically observed in traditional multi-task settings.

3.1. Preliminaries

Model merging. Given a collection of D fine-tuned model weights $\{\theta_1, \theta_2, \dots, \theta_D\}$ from the same pre-trained

model architecture θ_{pre} , we first compute the set of delta weights $\Delta w = \{\Delta w_d \mid \Delta w_d := \theta_d - \theta_{pre}\}_{d=1}^D$, *i.e.*, the parameter-wise difference between the fine-tuned and pre-trained model and then merge all delta weights using a merging function f_{merge} , such that $\Delta_{merged} = f_{merge}(\Delta w, \lambda)$. Here, f_{merge} represents a merging method from the literature (*e.g.*, Section 2), and $\lambda \in \mathbb{R}$ is a scaling factor. After that, we sum Δ_{merged} to the pre-trained model as $\theta_{final} = \theta_{pre} + \Delta_{merged}$.

Domain generalization. In this work, we consider a scheme similar to multi-source Domain Generalization (DG) setting [51, 52], where we are given D distinct source domains $\mathcal{S} = \{S^d\}_{d=1}^D$. A domain d is defined by a joint distribution P_{XY}^d over a domain-specific input space \mathcal{X}^d and a shared target space \mathcal{Y} . Each source domain S^d is associated with a unique distribution P_{XY}^d , such that $P_{XY}^d \neq P_{XY}^{d'}$ for $d \neq d'$. The objective is to learn a single predictive model $f : \mathcal{X} \rightarrow \mathcal{Y}$ using only the data from \mathcal{S} . This model must generalize to an unseen target domain T associated with a novel distribution P_{XY}^T , where $T \notin \mathcal{S}$. Typical approaches run a single training containing all source-domain data and rely on a domain-generalization strategy [36, 50] or specifically designed loss function [9, 20, 23].

Model merging for domain generalization. We introduce a novel perspective on domain generalization through model merging as a compositional mechanism that enables better generalization to unseen domains. Our study investigates whether merging models fine-tuned on data from different distribution shifts can result in a single model with improved generalization to previously unseen data. This setting mirrors practical use cases of generalization, such as combining distinct visual styles or object concepts in generative models, or enabling zero-shot transfer to new tasks.

We evaluate merging in a leave-one-domain-out protocol [3, 9, 52]: given D domains $\{S^d\}_{d=1}^D$ we hold out a domain S^t as the unseen target and use the remaining $D - 1$ domains as sources. For each source domain S^d we assume access only to a fine-tuned parameter checkpoint θ_d (obtained by fine-tuning a common pre-trained backbone θ_{pre} on S^d), while source training data is unavailable. We form the set of delta parameters $\Delta w = \{\Delta w_d\}_{d \neq t}$ with $\Delta w_d = \theta_d - \theta_{pre}$ and compute a merged delta as:

$$\Delta_{merged} = f_{merge}(\{\Delta w_d\}_{d \neq t}; \lambda),$$

where f_{merge} is a data-free and optimization-free model merging method. The hyper-parameter λ controls the contribution of the merged weights to the pre-trained model. Prior work often tunes λ using validation data [10, 14, 26, 42], but in our setting, accessing the out-of-distribution data is prohibited, so we fix $\lambda = 1$. We evaluate the resulting

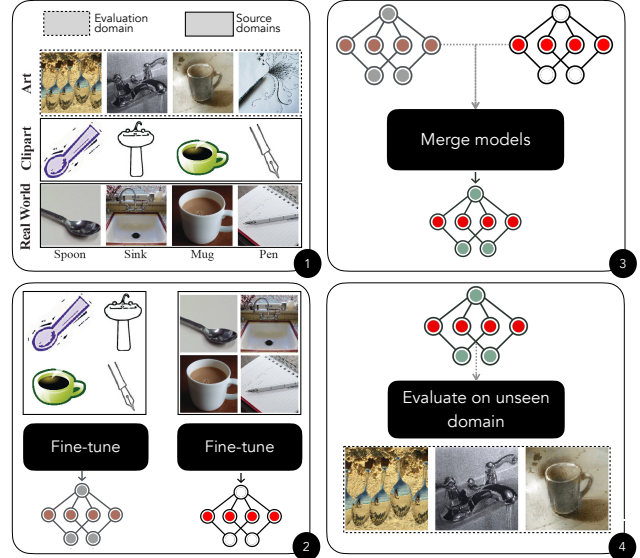


Figure 1. Evaluating model merging for domain generalization. We merge fine-tuned models across multiple domains and evaluate the merged model on an unseen domain. We repeat the leave-one-out process for each available domain.

model $\theta_{final} = \theta_{pre} + \Delta_{merged}$ on the test set of S^t . We repeat this procedure for each available domain $d = 1, \dots, D$. Figure 1 shows the protocol for a single evaluation domain.

3.2. Motivation

In traditional multi-task learning, the evaluation task is part of the set of models being merged, and the goal is to preserve information from all tasks so the merged model performs well across them. These tasks are diverse and different both in label space and in semantics, *e.g.* MNIST (digit classification) and RESIC45 (landscape classification).

On the other hand, domain generalization presents a more challenging scenario for model merging. Models trained on different domains often share the same label space but differ in their data distributions, encouraging each model to learn class-discriminative and domain-specific representations. Because these models solve similar classification problems, their delta weights tend to align along similar singular directions. This strong subspace overlap contrasts sharply with the diversity of multi-task learning. This overlap leads to interference among singular directions during merging, as competing dominant directions may translate into competing features in the merged model. Effectively handling this conflict becomes crucial for improving generalization in unseen domains.

To quantify the overlap between subspaces, we use the *Subspace Alignment Ratio* (SAR) metric [27]. The SAR metric between two delta weights Δw_i and Δw_j is defined as:

$$\text{SAR}(\Delta w_i, \Delta w_j; k_j) = \frac{\|\Pi_{k_j, j} \Delta w_i\|_F}{\|\Delta w_i\|_F}, \quad (1)$$

where $\Pi_{k_j, j} = U_{k_j, j} U_{k_j, j}^\top$ is the projection matrix onto the subspace spanned by the top k_j left-singular vectors of Δw_j . The columns of $U_{k_j, j}$ are obtained from the SVD decomposition of Δw_j , and the number of singular vectors k_j is determined from Δw_j by minimizing the approximation error ϵ :

$$k_j = \min \left\{ k : \|\Delta w_j - \Pi_{k_j, j} \Delta w_j\|_F \leq \epsilon \|\Delta w_j\|_F \right\} \\ = \min \left\{ k : \frac{\sum_{r=i=k+1}^r \sigma_i^2}{\sum_{i=1}^r \sigma_i^2} \leq \epsilon^2 \right\}, \quad (2)$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$ contains the singular values of Δw_j , and the equivalence follows from the definition of the Frobenius norm.

The SAR quantifies the alignment between the subspaces of two delta matrices as a function of the number of dominant singular vectors of Δw_j . We denote SAR_{avg} the *Average Subspace Alignment Ratio* across all layers to provide a single score measuring the overlap between two models.

Figure 2 shows the average SAR metric between tasks for merging 8 datasets in a multi-task setting [14], and between domains in domain generalization for the 6 domains in the DomainNet dataset. In multi-task learning, the overlap is more pronounced among groups of related tasks, such as MNIST, SVHN (digit classification), and EuroSAT, RESISC45 (landscape classification). In contrast, domain generalization produces consistently higher SAR values, even though all domains correspond to the same underlying task but differ in data distribution. We also report an additional overlap measure based on principal angles [19] in Appendix Section A2.4

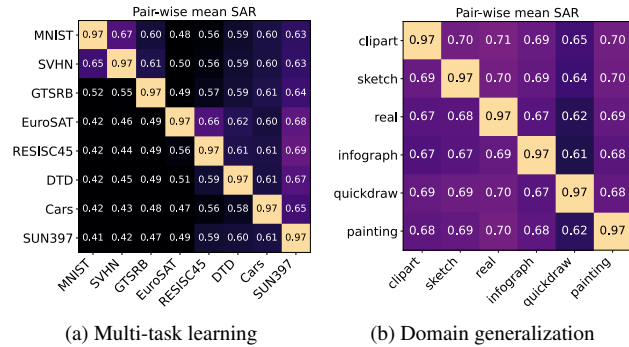


Figure 2. Pairwise similarity of fine-tuned models measured by the Subspace Alignment Ratio (SAR). (a) Models fine-tuned for 8 datasets from Ilharco et al. [14]. (b) Models fine-tuned for 6 domains of DomainNet [32]. The similarity between multi-domain models is much higher than between multi-task models, with much more overlap between the subspaces occupied, creating opportunity for conflicts.

We argue that domain shifts induce subspaces that overlap more strongly than task shifts. Such overlap introduces additional challenges for SVD-based merging methods. These methods prioritize the top singular directions and their corresponding eigenvalues. However, when fine-tuned models across different domains share similar singular directions, conflicts can arise during inference. In practice, these competing directions from one domain can dominate those from another, potentially biasing the merging process toward tasks with larger singular values and, thus, harming out-of-distribution generalization performance.

4. SCORE

In Section 3.2, we showed that there is substantial overlap between subspaces in domain generalization, which translates to higher conflicting singular directions and feature competition in the merged model. To address this issue, we propose SCORE, a Subspace CONflict-Resolving mERging method designed to alleviate such singular subspace conflicts. SCORE constructs a shared orthogonal basis by computing the principal components of the concatenated leading singular vectors of all models fine-tuned under domain shift. Each task matrix is then projected into this shared basis, where off-diagonal components shows the inter-domain conflicts in singular directions. We outline our method in Algorithm 1 for a single layer, and describe each step in the following.

Algorithm 1 SCORE

Require: Pre-trained model θ_{pre} , Task matrices $\{\Delta_1, \dots, \Delta_D\}$, number of domains to merge D

Ensure: Merged task matrix \hat{M}

- 1: **for** $d = 1$ to D **do**
 - 2: Compute SVD: $\Delta_d = U_d \Sigma_d V_d^\top$
 - 3: Retain first $\frac{1}{D}$ singular components of U_d, V_d
 - 4: **end for**
 - 5: **Concatenate** the matrices:

$$U_* \leftarrow [U_1 \mid U_2 \mid \dots \mid U_D], V_* \leftarrow [V_1 \mid V_2 \mid \dots \mid V_D]$$
 - 6: Compute SVD of U_* : $P_{U_*} \Sigma_{U_*} Q_{U_*}^\top$
 - 7: Compute SVD of V_* : $P_{V_*} \Sigma_{V_*} Q_{V_*}^\top$
 - 8: $U_\perp \leftarrow P_{U_*} Q_{U_*}^\top$ \triangleright Find orthogonal basis for U_*
 - 9: $V_\perp \leftarrow P_{V_*} Q_{V_*}^\top$ \triangleright Find orthogonal basis for V_*
 - 10: $\Sigma_{score} \leftarrow \mathbf{0}$
 - 11: **for** $d = 1$ to D **do**
 - 12: $\Delta'_d \leftarrow U_\perp^\top \Delta_d V_\perp$ \triangleright Change basis
 - 13: $\Sigma_{score} \leftarrow \Sigma_{score} + \text{trim}(\Delta'_d)$ \triangleright Eq. 4
 - 14: **end for**
 - 15: $\hat{M} \leftarrow U_\perp \Sigma_{score} V_\perp^\top$ \triangleright Reconstruct merged matrix
 - 16: **return** $\theta_{pre} + \hat{M}$
-

Combining domain-specific directions. Consider a collection of delta weights for the l -th network layer $\Delta w^{(l)} = \{\Delta w_d^{(l)}\}_{d=1}^D$ where each $\Delta w_d^{(l)} = \theta_d^{(l)} - \theta_{pre}^{(l)} \in \mathbb{R}^{m \times n}$ represents the parameter difference for domain d in layer l . Applying Singular Value Decomposition (SVD) to $\Delta w_d^{(l)} = U \Sigma V^T$, where $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$ are the left and right singular vectors, respectively, and $\Sigma \in \mathbb{R}^{r \times r}$ is a diagonal matrix containing the singular values. For clarity, we omit the layer superscript l in the following discussion and describe the merging process for a single layer.

To accommodate the singular directions of each individual domain, we compute the SVD decomposition of each delta matrix and consider only the top left and right singular directions:

$$U_* \leftarrow [U_1 | U_2 | \dots | U_D], V_* \leftarrow [V_1 | V_2 | \dots | V_D],$$

where $U_* \in \mathbb{R}^{m \times r}$ and $V_* \in \mathbb{R}^{n \times r}$ are the concatenated top- k domain-specific left and right singular directions from the SVD decomposition of Δw_d , respectively.

Orthogonalization. There is no guarantee that U_* and V_* are orthogonal to each other, which can compromise the reconstruction property of the SVD if used as they are. To create a shared subspace across domains that sticks to SVD constraints, we need to orthogonalize U_* and V_* . Following Gargiulo et al. [8], we compute the SVD of $U_* = P_{U_*} \Sigma_{U_*} Q_{U_*}^T$ and $V_* = P_{V_*} \Sigma_{V_*} Q_{V_*}^T$:

$$U_{\perp} \leftarrow P_{U_*} Q_{U_*}^{\top}, V_{\perp} \leftarrow P_{V_*} Q_{V_*}^{\top}. \quad (3)$$

This step yields U_{\perp} and V_{\perp} , which form a shared input and output basis that is closest to all D domain-specific subspaces.

Isolating subspace conflicts. Our intuition is that if the domain-specific top singular directions are mostly orthogonal to each other, both (U_{\perp}, V_{\perp}) defines a representative shared basis, with minimal level of conflicting dominant directions. However, as the subspace overlap is not negligible in domain generalization (see Section 3.2), we need to quantify which dominant directions are conflicting and their respective magnitude. For this, we apply the change of basis for each domain-specific matrix using U_{\perp} and V_{\perp} to understand which dominant directions each domain-specific weights Δ_d use:

$$\Delta'_d = U_{\perp}^{\top} \Delta_d V_{\perp},$$

where $\Delta'_d \in \mathbb{R}^{r \times r}$. The conflict in singular directions becomes explicit by reformulating the domain-specific SVD $\Delta_d = U_d \Sigma_d V_d$ with the shared projection:

$$\Delta'_d = U_{\perp}^{\top} (U_d \Sigma_d V_d^{\top}) V_{\perp} = (U_{\perp}^{\top} U_d) \Sigma_d (V_d^{\top} V_{\perp}).$$

Let us define $R_U^{(d)} = U_{\perp}^{\top} U_d$ and $R_V^{(d)} = V_d^{\top} V_{\perp}$. These matrices represent the dot products between the shared basis vectors and the domain d 's specific basis vectors. The

change of basis isolates the conflicts in dominant singular directions by transforming the matrix so that we can clearly distinguish between two types of information:

- 1. Agreement (Diagonal):** The diagonal elements $(\Delta'_d)_{ii}$ measure the magnitude of domain d along the i -th *shared* principal direction.
- 2. Conflict (Off-Diagonal):** The off-diagonal elements $(\Delta'_d)_{ij}$, for $i \neq j$ captures how domain d couples the i -th and j -th shared directions. Large off-diagonals indicate cross-talk between shared coordinates.

Figure 3 illustrates how block structures emerge after changing the basis for three distinct attention blocks in the ViT-B-32 model. For instance, we consider 6 DomainNet domains and then calculate $\Sigma_{score} = \sum_{d=1}^{D=6} U_{\perp}^{\top} \Delta_d V_{\perp}$. We present three scenarios, from left to right:

- **High agreement and minimal off-diagonal conflicts:** most energy lies on the diagonal, indicating the shared basis captures the dominant domain directions and domains align with the principal coordinates;
- **High agreement and moderate off-diagonal conflicts:** diagonal entries still dominate the magnitudes, but noticeable off-diagonal coupling appears, suggesting conflicts between shared directions;
- **Low agreement and moderate off-diagonal conflicts:** main diagonal energy is weak, and cross-terms are evident, indicating the shared basis failed to capture many domain-specific top directions.

Trimming off-diagonal outliers. Building on the previous analysis, we intend to keep the main diagonal elements of Δ'_d , which capture the agreement between each domain and the shared principal directions. However, keeping only the diagonal entries may miss important off-diagonal components that carry significant information about shared variability across domains. Therefore, instead of discarding all off-diagonal terms, we retain those with significant magnitude while suppressing outlier or noisy couplings.

Finally, to consolidate multiple Δ'_d and reconstruct the merged matrix as $\hat{M} = U_{\perp} \Sigma_{score} V_{\perp}^{\top}$, we compute the final directions $\Sigma_{score} = \sum_{d=1}^D \text{trim}(\Delta'_d)$, where the *trim* can be described as:

$$\text{trim}(\Delta'_k)_{ij} = \begin{cases} (\Delta'_k)_{ii}, & \text{if } i = j \quad (\text{Keep the diagonal}) \\ (\Delta'_k)_{ij}, & \text{if } i \neq j \text{ and } |(\Delta'_k)_{ij} - \mu_{\text{off}}| < \tau \cdot \sigma_{\text{off}} \\ 0, & \text{otherwise} \quad (\text{Prune outliers}), \end{cases} \quad (4)$$

where μ_{off} and σ_{off} are the mean and standard deviation of all off-diagonal elements. Considering only the main diagonal indicates taking account of dominant directions; on the other hand, keeping the whole Δ'_d account for all pair-wise interferences may introduce high levels of noise in the final and merged weights. We set $\tau = 1.96$ as it corresponds to 95% confidence interval in a standard normal distribution.

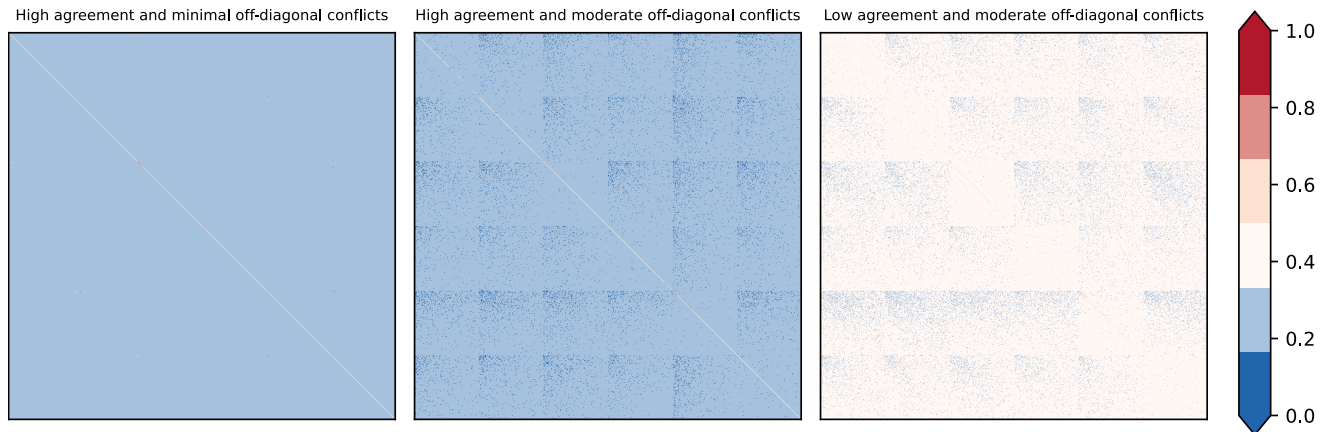


Figure 3. Σ_{score} (see Alg. 1) computed over 6 DomainNet domains across different attention layers of a ViT-B-32. Left: high agreement/low conflict, with a strong main diagonal and clean off-diagonal. Middle: high agreement/high conflict, with both the main diagonal and off-diagonal blocks significantly occupied. Right: low agreement/high conflict, with off-diagonal elements dominating and main diagonal nearly absent, indicating that the shared basis (U_{\perp} and V_{\perp}) fails to capture consistent singular directions across domains and that those domains intensely compete for subspaces in the shared representation.

5. Experiments

We evaluate three CLIP [34] variants (ViT-B/32, ViT-B/16, and ViT-L/14) on a comprehensive benchmark of eight domain generalization image classification datasets. The benchmark includes six natural-image datasets: PACS [21], DomainNet [32], ImageNet-R [11], NICOpp [51], OfficeHome [41], and TerraIncognita [2]; and two medical datasets, FedISIC [30] and RetinaDomains [6, 17, 22, 33]. In total, these datasets cover 49 domains with varying label spaces (4-365 classes) and capture a broad spectrum of distribution shifts. Details about each dataset appear in Appendix A1, and per-domain results appear in the Appendix A2.5.

We compare our approach with the following model merging methods: Task Arithmetic [14], TIES [45], MagMax [26], PCB [10], TSV [8], ISO-C [27], and ISO-CTS [27]. We report the absolute accuracy for natural-images datasets, and balanced accuracy for medical datasets, since they are highly imbalanced. For reference, we include each model’s zero-shot performance as a lower bound, and tasks experts as upper-bound. The code to reproduce our work is available at https://github.com/VirtualSpaceman/score_cvpr26.

Evaluation protocol. We adopt the leave-one-domain-out evaluation protocol [9, 52]. For each held-out target domain, we merge the fine-tuned models from all remaining source domains and evaluate the merged model on the unseen target domain, without any additional training or adaptation. We repeat this process for every domain and report the mean performance across all target domains.

Fine-tuning details. We fully fine-tune the CLIP’s image encoder with a batch size of 128, a learning rate of $1e-5$

coupled with a cosine annealing schedule, and AdamW optimizer [24] with weight decay 0.1 while keeping the text encoder unchanged. We use the CLIP’s text encoder as the final classification layer output and keep it frozen during fine-tuning. This fine-tuning recipe preserves the model’s open-vocabulary nature without compromising accuracy [8, 14, 26].

5.1. Domain generalization performance

Table 1 compares the performance across several model merging methods for ViT-B-32 and ViT-L-14. Results for ViT-B-16 appear in the appendix. Our approach achieves the highest average accuracy across all methods, surpassing the next-best competitor by 0.74 percentage points (p.p.) on the ViT-B-32 model and 0.58 p.p. on the ViT-L-14 model. For the ViT-B-32, our method outperforms all competitors on the DomainNet dataset by 0.15 p.p., on NICOpp by 0.40 p.p., and on OfficeHome by 0.49 p.p. The performance gains are even more pronounced with the larger ViT-L-14 model. Our approach achieves the best results on PACS, tying for first place with a 0.02 p.p. advantage over the second-best, and on DomainNet, where we are also tied for first, exceeding the second-best by 0.63 p.p. Additionally, we outperform competitors by 1.01 p.p. on NICOpp and by 1.18 p.p. on TerraIncognita, while remaining highly competitive on ImageNetR and OfficeHome.

When evaluating on specialized medical datasets, all methods show performance variations with model scale. With the ViT-B/32 model, our approach ranks second on the FedISIC dataset. Using the larger ViT-L/14 model, it again secures the second-best result on FedISIC and achieves the top performance on RetinaDomains, demonstrating strong adaptability across challenging medical imaging domains.

Method	PACS	DomainNet	ImageNetR	NICOpp	OfficeHome	TerraIncognita	FedISIC	RetinaDomains	Avg	
ViT-B-32	Zero-shot	94.85	56.10	66.91	77.14	79.82	22.73	23.08	55.63	
	Task. Arithm.	93.40	5.40	4.65	56.92	66.73	40.47	32.18	40.00	42.47
	TIES	95.77	53.22	72.98	78.18	83.27	47.03	50.42	37.84	64.84
	MagMax	95.46	43.87	68.38	76.18	80.37	43.56	43.32	41.06	61.53
	PCB	95.35	<u>58.53</u>	<u>74.67</u>	<u>80.19</u>	83.59	40.84	38.41	35.45	63.38
	TSV	95.83	55.64	74.23	79.78	83.74	47.68	44.22	38.44	64.95
	Iso-C	95.37	55.84	74.74	79.81	82.34	36.53	32.91	32.97	61.31
	Iso-CTS	95.30	56.80	73.46	79.76	82.00	32.87	29.63	33.61	60.43
	SCORE (Ours)	95.70	58.68	73.44	80.59	84.23	46.92	<u>50.12</u>	35.87	65.69
	Task Experts	99.81	74.43	86.70	84.34	91.65	82.15	61.26	48.72	78.63
ViT-L-14	Zero-shot	98.39	<u>66.25</u>	86.51	84.06	88.82	46.46	19.77	26.00	64.53
	Task. Arithm.	98.53	30.61	13.07	73.39	89.97	51.11	31.64	42.81	53.89
	TIES	98.60	65.80	90.16	<u>85.53</u>	91.26	<u>59.72</u>	47.48	41.11	72.46
	MagMax	98.62	62.24	88.15	84.82	91.41	55.16	51.69	<u>40.74</u>	71.60
	PCB	98.54	68.26	90.38	86.44	91.13	59.32	38.93	38.22	71.40
	TSV	98.56	66.79	90.83	86.14	91.22	60.01	43.09	40.46	72.14
	Iso-C	98.51	66.89	90.90	86.20	90.32	57.52	33.31	34.55	69.78
	Iso-CTS	98.50	67.63	90.46	86.21	89.94	55.99	28.71	36.26	69.21
	SCORE (Ours)	98.62	68.26	<u>90.56</u>	86.54	<u>91.28</u>	61.19	<u>46.77</u>	41.12	73.04
	Task Experts	99.92	81.62	91.58	88.10	96.15	88.61	67.88	60.84	84.34

Table 1. Results for the leave-one-domain-out protocol: we measure accuracy (in %) on the left-out domain of each column. On average, our SCORE outperforms all existing solutions. Columnwise, best results in **bold** and second-best underlined: SCORE beats or stays very close to the best previous method.

5.2. Analysis and ablations

Model merging can outperform zero-shot evaluation on unseen domains. We investigate whether model merging can serve as a compositional mechanism to integrate knowledge from models fine-tuned on different distribution shifts. Specifically, we ask whether any merging strategy can surpass the zero-shot baseline in out-of-distribution generalization. Intuitively, combining models trained on complementary shifts should yield a merged model that generalizes better to unseen domains. However, this outcome is not guaranteed. If the compositional knowledge captured by the merged weights does not align with features that promote generalization, the merging process may instead degrade performance. Our study empirically tests this hypothesis, examining if model merging succeeds to outperform zero-shot evaluation in unseen environments.

Table 1 indicates that most strategies outperform the zero-shot evaluation on out-of-distribution, with Task Arithmetic being the lone exception. The latter was one of the first strategies beyond the simple average that relies on delta weights, but it is sensitive to the choice of hyperparameter. Beyond these methods, recent methods, including ours, significantly outperform this baseline. These results support our initial hypothesis that practitioners can rely on model merging to combine multiple sources of information into a single model with improved generalization performance.

These results indicate that model merging provides a practical and efficient strategy for combining multiple fine-tuned models to improve generalization under distribution shifts. This capability is valuable in dynamic or data-limited environments, where practitioners have access only to a fine-tuned model for certain distribution shifts, available

through model repositories. Performance gains are particularly evident on medical datasets. Model merging strategies consistently improved over the zero-shot baseline across all model sizes by at least 6 p.p. (FedISIC) and 8 p.p. (RetinaDomains) for ViT-B-32, comparing the baseline with the worst model merging strategy on that particular dataset.

Model merging outperforms model ensembles in domain generalization. We compare model merging against traditional model ensembling to understand their relative effectiveness and efficiency. Model merging operates directly in parameter space, producing a single merged model that carries the combined knowledge of multiple fine-tuned models without increasing inference cost. In contrast, model ensembling combines predictions in output space, which scales poorly as the number of source domains increases. Each additional model in an ensemble adds memory overhead, computational cost, and inference latency, since every model must be loaded in memory, and its outputs aggregated at test time.

Table 2 compares the best-performing model-merging method according to our experiments with the logit ensemble, a common ensemble strategy [7, 48] on the same 8 domain generalization datasets. Our SCORE method consistently outperforms the logit ensemble baseline by 1.12 to 1.90 p.p. across all model sizes. Critically, model merging maintains the inference cost of a single model, thereby alleviating the computational and memory overhead of an ensemble’s multiple forward passes. Surpassing the ensemble’s performance, rather than merely matching it, demonstrates that SCORE successfully integrates complementary knowledge from the fine-tuned models.

Table 2. Mean average accuracy (%) across datasets for different ViT architectures. Model ensemble is the logit ensemble baseline; values in parentheses denote relative changes. Improvements are shown in green.

Method	ViT-B-32	ViT-B-16	ViT-L-14
Model ensemble	<u>64.57</u>	<u>68.07</u>	<u>71.81</u>
SCORE (Ours)	65.69 (+1.12)	69.97 (+1.90)	73.04 (+1.24)

Our findings expand the previous evaluation by Wortsman et al. [43], which considered only weight averaging and distribution shifts on ImageNet. In this study, we consider a broader range of datasets for domain generalization and model-merging candidates, using a typical evaluation scheme from the domain generalization literature.

The impact of the trim function on performance.

SCORE employs a *trim* operation to mitigate conflicts among singular directions and enhance generalization. It retains the main diagonal while removing off-diagonal outliers from Σ_{score} to reduce conflicting directions. Table 3 shows our ablation study to considering different strategies for Σ_{score} :

- **Diagonal only:** keeping only the main diagonal elements;
- **Off-diagonal only:** considering only the off-diagonal components (keep only the domain-wise conflicts);
- **Full matrix:** preserving both diagonal and off-diagonal elements;
- **Trimmed:** retaining both diagonal and off-diagonal elements but trimming the off-diagonal outliers (*trim* operation). Our proposed SCORE relies on this strategy.

Table 3. Ablation on the choice of merged task matrix. We report the mean average accuracy (%) across domains for each ViT architecture and, in parenthesis, the positive or negative difference to the baseline given by Diagonal. The proposed trimmed merging offers the best results.

Method	ViT-B-32	ViT-B-16	ViT-L-14
Diagonal	<u>63.62</u>	<u>67.38</u>	<u>71.50</u>
Off-diagonal	58.41 (-5.21)	62.41 (-4.97)	67.46 (-4.04)
Full matrix	7.59 (-56.03)	7.66 (-59.72)	7.70 (-63.80)
Trimmed	65.69 (+2.07)	69.97 (+2.59)	73.04 (+1.53)

We use the configuration that keeps only the main diagonal elements of Σ_{score} as our baseline. Retaining only the off-diagonal components reduces average performance across all architectures by up to 5 p.p., indicating that these components still carry meaningful information for generalization. However, preserving both diagonal and off-diagonal elements without trimming leads to a substantial performance drop, suggesting strong interference and competition between singular directions, promoting high feature

competition in the merged model. In contrast, our proposed strategy via the *trim* function balances the contributions of shared subspaces while removing destructive outliers, improving by up to 2.59 p.p. in average performance relative to the diagonal-only baseline.

6. Conclusion

We studied model merging as a compositional mechanism for integrating multiple fine-tuned models trained on distinct domains and evaluated the merged models’ performance on unseen domains. We introduced SCORE, a novel model merging approach to mitigate conflicts over singular directions. Our method outperforms, on average, all other model merging methods for 8 domain generalization datasets and 3 model sizes. Also, SCORE outperforms a model-ensembling baseline. Our proposed contribution is agnostic to the specifics of the domain generalization method and depends only on practitioners’ fine-tuning from a common base model. Model merging is complementary to most domain generalization methods. Future research could investigate the benefits of combining multiple models, each fine-tuned on the same multi-source domain data but utilizing different domain generalization losses. Our work also paves the way for using model merging as a strategy to explore generalization capabilities, similar to the compositional experiments conducted by Kempf et al. [18], Mayilvahanan et al. [29]. Rather than retraining models from scratch, model merging allows practitioners to merge individual models and assess whether the compositional results from Kempf et al. [18] still hold.

Limitations. Our experimental design assumes we have access only to the fine-tuned model for a given distribution shift, not to the source data. A limitation of every parameter-wise model merging is only merging models that share the same architecture and fine-tune from the same seed model. While our work primarily focuses on image classification, we anticipate future extensions to natural language processing and generative models, including large language models and image generative models.

Acknowledgements LC is partially funded by FAPESP (2024/16685-7), CAPES, and Becas Santander/UNICAMP 2022. SA is also partially funded by FAPESP (2023/12086-9, 2023/12865-8, 2020/09838-0, 2013/08293-7), H.IAAC 01245.003479/2024-10 and CNPq 316489/2023-9. CZ and RB acknowledge the Gauss Centre for Supercomputing e.V. for funding this project by providing computing time on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC). They also acknowledge funding from the European Research Council (ERC) under the Horizon Europe Framework Programme (HORIZON) for proposal number 101116395 SPARSE-ML.

References

- [1] Samuel Ainsworth, Jonathan Hayase, and Siddhartha Srinivasa. Git re-basin: Merging models modulo permutation symmetries. In *International Conference on Learning Representations*, 2023. 2
- [2] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *European Conference on Computer Vision*, 2018. 6, 1
- [3] Levy Chaves, Eduardo Valle, and Sandra Avila. Weight weaving: Parameter pooling for data-free model merging. In *Advances in Neural Information Processing Systems Workshops Unifying Representations in Neural Models*, 2025. 3
- [4] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kalloo, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (ISIC). In *International Symposium on Biomedical Imaging*, 2018. 1
- [5] MohammadReza Davari and Eugene Belilovsky. Model breadcrumbs: Scaling multi-task model merging with sparse masks. In *European Conference on Computer Vision*, 2024. 2
- [6] Etienne Decencière, Xiwei Zhang, Guy Cazuguel, Bruno Lay, Béatrice Cochener, Caroline Trone, Philippe Gain, John-Richard Ordóñez-Varela, Pascale Massin, Ali Erginay, et al. Feedback on a publicly distributed image database: the messidor database. *Image Analysis & Stereology*, 2014. 6, 1
- [7] Mudasir A Ganaie, Minghui Hu, Ashwani Kumar Malik, Muhammad Tanveer, and Ponnuthurai N Suganthan. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence*, 2022. 7
- [8] Antonio Andrea Gargiulo, Donato Crisostomi, Maria Sofia Bucarelli, Simone Scardapane, Fabrizio Silvestri, and Emanuele Rodola. Task singular vectors: Reducing task interference in model merging. In *Conference on Computer Vision and Pattern Recognition Conference*, 2025. 1, 2, 5, 6
- [9] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. 3, 6
- [10] Du Guodong, Junlin Lee, Jing Li, Runhua Jiang, Yifei Guo, Shuyang Yu, Hanting Liu, Sim Kuan Goh, Ho-Kin Tang, Daojing He, et al. Parameter competition balancing for model merging. In *Neural Information Processing Systems*, 2024. 2, 3, 6
- [11] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. *International Conference on Computer Vision*, 2021. 6, 1
- [12] Carlos Hernández-Pérez, Marc Combalia, Sebastian Podlipnik, Noel CF Codella, Veronica Rotemberg, Allan C Halpern, Ofer Reiter, Cristina Carrera, Alicia Barreiro, Brian Helba, et al. Bcn20000: Dermoscopic lesions in the wild. *Scientific data*, 2024. 1
- [13] Chenyu Huang, Peng Ye, Tao Chen, Tong He, Xiangyu Yue, and Wanli Ouyang. Emr-merging: Tuning-free high-performance model merging. *Advances in Neural Information Processing Systems*, 2024. 2
- [14] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *International Conference on Learning Representations*, 2023. 1, 2, 3, 4, 6
- [15] Xisen Jin, Xiang Ren, Daniel Preotiu-Pietro, and Pengxiang Cheng. Dataless knowledge fusion by merging weights of language models. In *International Conference on Learning Representations*, 2023. 2
- [16] Keller Jordan, Hanie Sedghi, Olga Saukh, Rahim Entezari, and Behnam Neyshabur. Repair: Renormalizing permuted activations for interpolation repair. In *International Conference on Learning Representations*, 2023. 2
- [17] Karthik, Maggie, and Sohier Dane. Aptos 2019 blindness detection. <https://kaggle.com/competitions/aptos2019-blindness-detection>, 2019. Kaggle. 6, 1
- [18] Elias Kempf, Simon Schrod, Max Argus, and Thomas Brox. When and how does clip enable domain and compositional generalization? In *International Conference on Machine Learning*, 2025. 8
- [19] Andrew V. Knyazev and Merico E. Argentati. Principal angles between subspaces in an a-based scalar product: Algorithms and perturbation estimates. *SIAM Journal on Scientific Computing*, 23(6):2008–2040, 2002. 4, 1
- [20] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, 2021. 3
- [21] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *International Conference on Computer Vision*, 2017. 6, 1
- [22] Tao Li, Yingqi Gao, Kai Wang, Song Guo, Hanruo Liu, and Hong Kang. Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Information Sciences*, 501:511 – 522, 2019. 6, 1
- [23] Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, 133(1), 2025. 3
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 6, 1
- [25] Zhenyi Lu, Chenghao Fan, Wei Wei, Xiaoye Qu, Danyang Chen, and Yu Cheng. Twin-merging: Dynamic integration of modular expertise in model merging. *Advances in Neural Information Processing Systems*, 2024. 2
- [26] Daniel Marczak, Bartłomiej Twardowski, Tomasz Trzcíński, and Sebastian Cygert. Magmax: Leveraging model merging for seamless continual learning. In *European Conference on Computer Vision*, 2024. 2, 3, 6
- [27] Daniel Marczak, Simone Magistri, Sebastian Cygert, Bartłomiej Twardowski, Andrew D Bagdanov, and Joost

- van de Weijer. No task left behind: Isotropic model merging with common and task-specific subspaces. In *International Conference on Machine Learning*, 2025. 2, 3, 6, 1
- [28] Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 2022. 2
- [29] Prasanna Mayilvahanan, Thaddäus Wiedemer, Evgenia Rusak, Matthias Bethge, and Wieland Brendel. Does clip’s generalization performance mainly stem from high train-test similarity? In *International Conference on Learning Representations*, 2024. 8
- [30] Jean Ogier du Terrail, Samy-Safwan Ayed, Edwige Cyffers, Felix Grimmer, Chaoyang He, Regis Loeb, Paul Mangold, Tanguy Marchand, Othmane Marfoq, Erum Mushtaq, et al. Flamby: Datasets and benchmarks for cross-silo federated learning in realistic healthcare settings. *Advances in Neural Information Processing Systems*, 2022. 6, 1
- [31] Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *Advances in Neural Information Processing Systems*, 2023. 2
- [32] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *International Conference on Computer Vision*, 2019. 4, 6, 1
- [33] Prasanna Porwal, Samiksha Pachade, Ravi Kamble, Manesh Kokare, Girish Deshmukh, Vivek Sahasrabudhe, and Fabrice Meriaudeau. Indian diabetic retinopathy image dataset (idrid), 2018. 6, 1
- [34] 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*, 2021. 6
- [35] Alexandre Ramé, Kartik Ahuja, Jianyu Zhang, Matthieu Cord, Léon Bottou, and David Lopez-Paz. Model ratatouille: Recycling diverse models for out-of-distribution generalization. In *International Conference on Machine Learning*, 2023. 2
- [36] Alexander Robey, George J Pappas, and Hamed Hassani. Model-based domain generalization. *Advances in Neural Information Processing Systems*, 2021. 3
- [37] George Stoica, Daniel Bolya, Jakob Brandt Bjorner, Pratik Ramesh, Taylor Hearn, and Judy Hoffman. Zipit! merging models from different tasks without training. In *International Conference on Learning Representations*, 2024. 2
- [38] George Stoica, Pratik Ramesh, Boglarka Ecsedi, Leshem Choshen, and Judy Hoffman. Model merging with svd to tie the knots. In *International Conference on Learning Representations*, 2025. 2
- [39] Derek Tam, Yash Kant, Brian Lester, Igor Gilitschenski, and Colin Raffel. Realistic evaluation of model merging for compositional generalization. *arXiv preprint arXiv:2409.18314*, 2024. 2
- [40] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 2018. 1
- [41] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Conference on Computer Vision and Pattern Recognition*, 2017. 6, 1
- [42] Ke Wang, Nikolaos Dimitriadis, Guillermo Ortiz-Jiménez, François Fleuret, and Pascal Frossard. Localizing task information for improved model merging and compression. In *International Conference on Machine Learning*, 2024. 2, 3
- [43] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, 2022. 2, 8
- [44] Zhengqi Xu, Ke Yuan, Huiqiong Wang, Yong Wang, Mingli Song, and Jie Song. Training-free pretrained model merging. In *Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [45] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 2023. 1, 2, 6
- [46] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. *International Conference on Learning Representations*, 2024. 2
- [47] Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang, Xiaochun Cao, Jie Zhang, and Dacheng Tao. Model merging in llms, mllms, and beyond: Methods, theories, applications, and opportunities. *ACM Computing Surveys*, 2026. 1
- [48] Yongquan Yang, Haijun Lv, and Ning Chen. A survey on ensemble learning under the era of deep learning. *Artificial Intelligence Review*, 2023. 7
- [49] Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *International Conference on Machine Learning*, 2024. 1, 2
- [50] Hanlin Zhang, Yi-Fan Zhang, Weiyang Liu, Adrian Weller, Bernhard Schölkopf, and Eric P Xing. Towards principled disentanglement for domain generalization. In *Conference on Computer Vision and Pattern Recognition*, 2022. 3
- [51] Xingxuan Zhang, Yue He, Renzhe Xu, Han Yu, Zheyang Shen, and Peng Cui. Nico++: Towards better benchmarking for domain generalization. In *Conference on Computer Vision and Pattern Recognition*, 2023. 3, 6, 1
- [52] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4), 2022. 3, 6