

---

# EMR-MERGING: Tuning-Free High-Performance Model Merging

---

Chenyu Huang<sup>1†</sup>, Peng Ye<sup>1,3†‡</sup>, Tao Chen<sup>1\*</sup>, Tong He<sup>2</sup>, Xiangyu Yue<sup>3</sup>, Wanli Ouyang<sup>3</sup>

<sup>1</sup> Fudan University    <sup>2</sup> Shanghai AI Laboratory

<sup>3</sup> The Chinese University of Hong Kong

cyhuang24@m.fudan.edu.cn

## Abstract

The success of pretrain-finetune paradigm brings about the release of numerous model weights. In this case, merging models finetuned on different tasks to enable a single model with multi-task capabilities is gaining increasing attention for its practicability. Existing model merging methods usually suffer from (1) significant performance degradation or (2) requiring tuning by additional data or training. In this paper, we rethink and analyze the existing model merging paradigm. We discover that using a single model’s weights can hardly simulate all the models’ performance. To tackle this issue, we propose ELECT, MASK & RESCALE-MERGING (EMR-MERGING). We first (a) elect a unified model from all the model weights and then (b) generate extremely light-weight task-specific modulators, including masks and rescalers, to align the direction and magnitude between the unified model and each specific model, respectively. EMR-MERGING is tuning-free, thus requiring no data availability or any additional training while showing impressive performance. We find that EMR-MERGING shows outstanding performance compared to existing merging methods under different classical and newly-established settings, including merging different numbers of vision models (up to 30), NLP models, PEFT models, and multi-modal models.<sup>1</sup>

## 1 Introduction

With the rapid development of deep learning, different model architectures [36, 22, 71, 88] are proposed, along with multiple training strategies [89, 86]. Pre-trained models’ capabilities are enhanced, thus showing increasing significance [54, 22, 7, 19]. Finetuning models on downstream tasks from a pre-trained model has become a standard paradigm in both NLP and vision fields [20, 51, 19, 22, 5, 87], which usually leads to improved performance with less labeled data. With the development of open-source repositories such as Huggingface [79], timm [77], and torchvision [44], the number of pre-trained and finetuned checkpoints exponentially rise. However, applying individual models to different tasks results in high storage and deployment costs. Multi-task learning (MTL) partially solves this problem by jointly training a model using multiple datasets [70, 93, 95], but it suffers from (i) high computational costs and (ii) data unavailability due to privacy [33]. Recently, model merging attempts to solve these drawbacks by combining weights instead of additional training, thus showing vital significance and broad application prospects.

A simple strategy of model merging is averaging the model weights [80], but it usually causes obvious performance degradation, as shown in Fig. 1. To this end, there are multiple model merging methods proposed to improve the performance of the merged model, which can be roughly divided into three

---

\*Corresponding Author(eetchen@fudan.edu.cn). †Equal Contribution. ‡Project Lead.

<sup>1</sup>Our code is available at [https://github.com/harveyhuang18/EMR\\_Merging](https://github.com/harveyhuang18/EMR_Merging).

categories: (i) *Weighted averaging of model weights* include Fisher-Merging [46] and RegMean [33]. They use pre-computed Fisher information matrices [23] and inner-product matrices [33] to tune the coefficients for weighted averaging. (ii) *Task vector-based methods* that add task vectors together instead of model weights, include Task Arithmetic [30], Ties-Merging [84], and AdaMerging [85]. Ties-Merging handles the interference issue and AdaMerging adaptively tunes the merging coefficients. (iii) *Pre-processing techniques* include DARE [90]. It reduces interference by dropping most elements and rescaling the others in task vectors. Despite the promising results, there are two unresolved problems with the existing model merging methods: (1) The performance gap between the merged model and individual models or MTL is still obvious, as shown in Fig. 1. (2) The performance improvement of existing methods depends on tuning by data or training, as shown in Tab. 1.

To boost the performance of model merging, we rethink and analyze the existing model merging paradigm. We discover that the goal of all the existing methods is to obtain a single model applicable to all the  $N$  tasks, as follows:

$$W_M = \mathcal{M}([W_1..W_N]), \quad (1)$$

where  $[W_1..W_N]$  are the model weights to be merged,  $\mathcal{M}$  denotes the merging function, and  $W_M$  is the merged model weight. This paradigm may inevitably lead to a non-negligible gap between the merged model and each individual model, especially when there are numerous models or models on challenging tasks. We argue that using a single model weight to simulate all the model weights is sub-optimal. To tackle this issue, we propose a brand new merging paradigm: We first extract a unified model weight from all the models’ weights, and then we calculate and store significant but lightweight task-specific parts of each model weight. This process can be written as:

$$W_{uni}, [E_1..E_N] = \mathcal{M}'([W_1..W_N]), \quad (2)$$

where  $W_{uni}$  represents the common and shared part of all model weights and  $[E_1..E_N]$  denote the task-specific parts of each model weight.  $\mathcal{M}'$  is the revised merging function following our paradigm.

Based on the above paradigm, we propose EMR-MERGING (ELECT, MASK & RESCALE-MERGING). We first elect a unified model from all the model weights. The election strategy is choosing the maximum absolute value of each parameter on the specified sign direction to minimize interference and avoid additional tuning. Then we generate additional lightweight task-specific modulators, including masks and rescalers. Their functions are respectively to align the direction and magnitude of the unified model with the original task-specific model. We find that applying the task-specific modulators to the unified model can better approximate the task-specific model, thus improving performance. The detailed process, theoretical and empirical analysis of the proposed method are illustrated in Section 3. By applying our method, the performance of model merging is significantly enhanced and is comparable to MTL or individual models, as shown in Fig. 1. Meanwhile, EMR-MERGING requires no data, tuning, or any additional training, as shown in Tab. 1.

We first demonstrate the effectiveness of the proposed EMR-MERGING under the existing setting of (1) merging Vision Transformer (ViT) [22] models of different sizes on 8 vision tasks, (2) merging parameter-efficient finetuning (PEFT) models on 11 language tasks, and (3) merging GPT-2 [55] models on 7 language tasks. Our method shows significant performance improvement under these settings, even when compared to the strongest baseline. We further validate the method’s effectiveness under newly-established and more challenging settings including: (4) merging ViTs on **30** vision




Figure 1: The average accuracy of the multi-task performance of different model merging methods on eight vision tasks. Among all the merging methods, our EMR-MERGING is the only one comparable to the performance of MTL and even individual models.

Table 1: Prerequisites for each method’s working.

Methods	Training-Data Tuning	Valid-Data inputs	Tuning labels	Tuning by Training
Weight Averaging	×	×	×	×
Traditional MTL	✓	×	×	✓
Fisher-Merging [46]	×	✓	×	×
RegMean [33]	×	✓	×	×
Task Arithmetic [30]	×	✓	✓	×
Ties-Merging [84]	×	✓	✓	×
AdaMerging [85]	×	✓	×	✓
<b>EMR-Merging(Ours)</b>	×	×	×	×

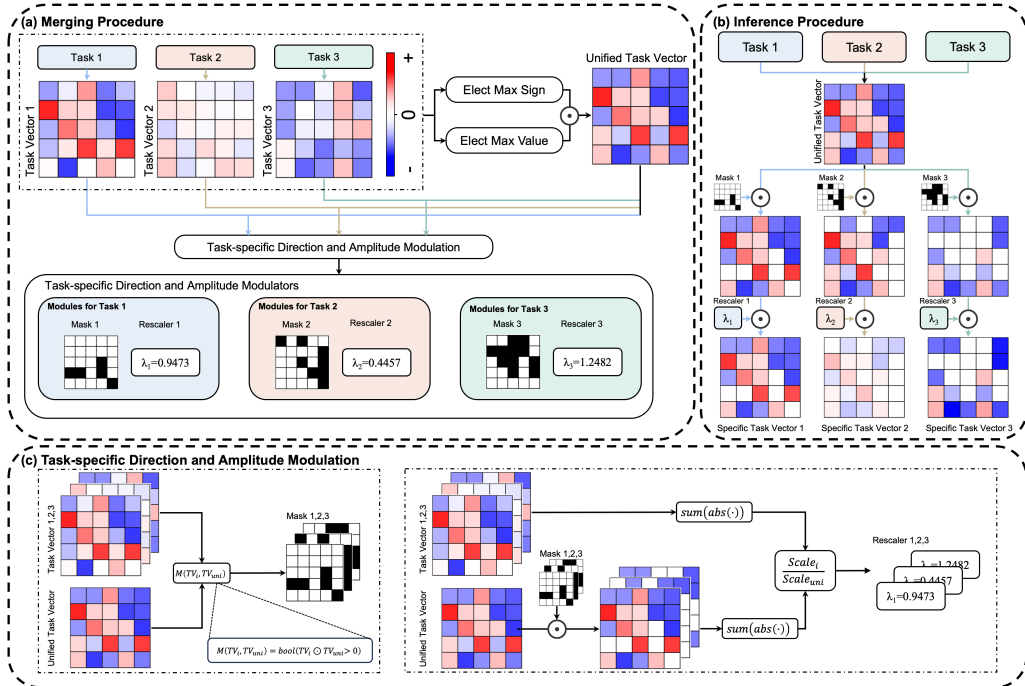


Figure 2: **Framework overview.** In the (a) Merging Procedure, we merge task-specific vectors into a unified task vector and lightweight task-specific modulators to modulate direction and amplitude. During the (b) Inference Procedure, we apply the corresponding mask and rescaler to the unified task vector to obtain a specific task vector. The process of (c) Task-specific Direction and Amplitude Modulation includes obtaining task-specific masks and scalars.

tasks, (5) merging RoBERTa [43] models on 8 NLP tasks, and (6) merging BEiT3 [75] models on 5 multi-modal tasks.

Our contributions can be summarized as: (1) We propose a novel merging method called EMR-MERGING, which merges task-specific models into a unified model and lightweight task-specific modulators (i.e., masks and rescalers), requiring no data, tuning, or additional training. (2) The proposed EMR-MERGING is simple-but-effective, and its effectiveness is validated on various classical benchmarks and newly-established benchmarks under various vision, NLP, PEFT, and multi-modal settings. (3) We show that the masks and rescalers of EMR-MERGING for aligning task-specific direction and amplitude of task vectors are applicable to most kinds of merging methods.

## 2 Related Work

**Model Merging** obtains a model using the existing task-specific model weights instead of training [33, 30, 84, 85, 66, 90, 46]. Simply averaging [80] usually causes severe performance degradation. Various methods are proposed to handle this problem. Fisher-Merging [46] and RegMean [33] use fisher information matrices [23] and inner-product matrices [33] to calculate the merging coefficients for weighted merging. However, they require additional matrices released by model owners or manually computed. Task Arithmetic [30] merges models by adding together task vectors, which is the difference between the finetuned and pre-trained models. Ties-Merging [84] and AdaMerging [85] are based on task vectors. Ties-Merging resolves interference and AdaMerging adaptively learns the merging coefficients. However, the performance of Task Arithmetic and Ties-Merging highly depends on manually tuning the merging coefficients and AdaMerging needs additional training to obtain them. DARE [90] reduces interference by randomly dropping most elements and rescaling the remaining ones in each task vector before merging. However, DARE’s performance is only validated under the setting of merging a limited number of tasks and the performance gain is also limited. In addition, all the existing methods merge models into a single one, and have not been verified under experimental settings of more models to merge, models on more difficult tasks, and

multi-modal models. In this paper, we propose EMR-MERGING, which requires no tuning while showing impressive performance under various settings.

**Multi-Task Learning** trains a single model using training data from multiple tasks together [70, 93, 95]. MTL typically necessitates access to the labeled data of multiple tasks for training the model from scratch. Though enabling the model multi-task capabilities, MTL suffers from not only (i) the expensive computational cost for training, especially for large models, but also (ii) the limited data availability due to data privacy [85]. In comparison, model merging solves the mentioned problems by combining the model weights without using training data or additional training, thus obtaining a multi-task model while sharply reducing the costs.

**Supervised Finetuning** from pre-trained models on down-stream tasks is becoming a standard paradigm in both NLP and vision fields [20, 51, 19, 22, 5]. Depending on whether all the parameters of models are adjusted, SFT can be divided into conventional full finetuning (FFT) and parameter-efficient finetuning (PEFT), which is proposed to reduce the number of trainable parameters for downstream tasks by adjusting the inserted small modules called adapters while keeping the whole model frozen [28, 29, 42]. PEFT is becoming the prevailing method to adapt pre-trained large models because of its efficiency [94]. There are a large number of pre-trained, full finetuned model weights, and PEFT module weights available on public repositories [79, 77, 44]. In this paper, the proposed EMR-MERGING is based on the common pretrain-finetune paradigm and we show the applicability of our method to both full finetuned models and PEFT modules.

### 3 Method

#### 3.1 Motivation

Given  $N$  tasks  $[T_1..T_N]$ , the goal of model merging is to obtain a model applicable to all the tasks using finetuned models  $[W_1..W_N]$  from the same pre-trained model  $W_{pre}$  on each task. Existing methods focus on merging the models into a single model  $W_M$ . Please check Appendix C for detailed information on the existing merging methods. However, a single model can hardly represent all the model weights, thus causing severe performance drops. We discover that the combination of a unified task vector and lightweight task-specific modulators can settle this issue to a significant extent by approximating the task-specific vectors better without any additional tuning. The size of proposed task-specific modulators is discussed in Section 4.4, which is much smaller than that of a model.

#### 3.2 ELECT, MASK & RESCALE-MERGING

The overall framework of EMR-MERGING is shown in Fig. 2. We follow the setting of task vector-based methods [30, 84, 85] and we merge models using task vectors. For task  $T_i$ ,  $i \in [1..N]$ , the corresponding task vector is defined as  $\tau_i = W_i - W_{pre}$ , where  $\tau_i \in \mathbb{R}^d$ .

**Electing a unified task vector** We first create an aggregate elected sign vector  $\gamma_{uni} = sgn(\sum_{t=1}^N \tau_t)$  by choosing the sign with the higher total magnitude of each parameter across all relevant task vectors. Then we choose the maximum absolute value of each parameter with the sign consistent with  $\gamma_{uni}$  from all the task vectors and obtain absolute value vector  $\epsilon_{uni} \in \mathbb{R}^d$ . By combining  $\gamma_{uni}$  and  $\epsilon_{uni}$ , the unified task vector can be obtained by  $\tau_{uni} = \gamma_{uni} \odot \epsilon_{uni}$ . The electing procedure can reserve the maximum amplitude and sign information shared by the task vectors, thereby maximally reducing interference. The unified task vector  $\tau_{uni}$  corresponds the  $W_{uni}$  in Eq. 2. Before being applied to task  $T_i$ , the  $\tau_{uni}$  needs to be modulated in advance by task-specific modulators, which are corresponding to  $E_i$  in Eq. 2. The generation of task-specific modulators is described below:

**Task-specific masks.** Next, we compare the unified task vector  $\tau_{uni}$  with each task vector  $\tau_i$ . The task-specific mask  $M_i = (\tau_i \odot \tau_{uni} > 0)$  for task  $i$  sets the elements whose signs are not correspondent with  $\tau_{uni}$  to zero and the rest to one. The function of the masks is to align the direction of the unified model with the task-specific model. The masks share the same structure with the task-specific models but due to their 1-bit nature, the size of a mask is much smaller than that of a task vector.

**Task-specific Rescalers** Then, for each task, we compute a rescaler parameter to keep the average absolute value of the elements in  $\tau_t$  and  $M_t \odot \tau_{uni}$  equal. The function of the rescalers  $\lambda_i = \frac{sum(abs(\tau_i))}{sum(abs(M_i \odot \tau_{uni}))}$  is to align the parameter magnitude of the unified model with the task-specific

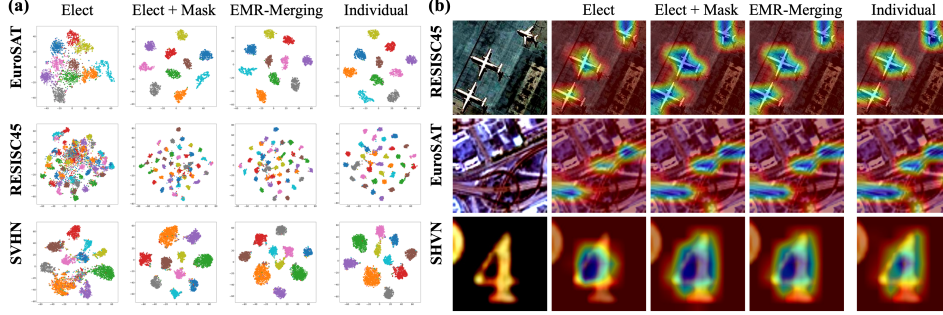


Figure 3: Partial (a) t-SNE and (b) Grad-CAM visualization results of EMR-MERGING’s procedures.

model. The significance of rescaling is also reported by DARE [90], which claims that after dropping most elements in a task vector, rescaling the rest leads to better results compared to not.

Before being applied to a task, a task-specific modulation is required to be conducted to the unified task vector. After that, we add it to the pre-trained parameter values  $W_{pre}$ . The inference steps of applying the merged model to task  $t$  are as follows:  $\hat{W}_t = W_{pre} + \hat{\tau}_t$ , where  $\hat{\tau}_t = \lambda_t \cdot M_t \odot \tau_{uni}$ . It should be noted that during the whole process, no additional tuning is needed, thus requiring no data or additional training. We summarize the algorithm flow in Appendix A.

### 3.3 Theoretical analysis

Our goal is to merge model weights by minimizing the distance between the merged model  $W_{uni}$  and each individual model  $W_i$ , where the distance can be calculated by:

$$Dis = \frac{\sum_{i=1}^N \|W_i - W_{uni}\|^2}{N} = \frac{\sum_{i=1}^N \|\tau_i - \tau_{uni}\|^2}{N} \quad (3)$$

where  $\tau_i$  refers to the task vector for task  $T_i$  and  $\tau_{uni}$  is the unified task vector.

**Analysis 1: Effectiveness of Masks.** After applying the masks  $M_i = (\tau_i \odot \tau_{uni} > 0)$  to the unified model  $\tau_{uni}$ , the distance  $Dis^M$  can be formulated as:

$$Dis^M = \frac{\sum_{i=1}^N \|\tau_i - M_i \odot \tau_{uni}\|^2}{N} \leq Dis \quad (4)$$

where  $Dis$  refers to the distance before applying the masks. Eq. 4 demonstrates that the distance between the merged model and each individual model can be reduced after applying the masks.

**Analysis 2: Effectiveness of Rescalers.** After applying the rescalers  $\lambda_i = \frac{sum(abs(\tau_i))}{sum(abs(M_i \odot \tau_{uni}))}$  to the masked task vectors  $M_i \cdot \tau_{uni}$ , the distance  $Dis^{M,\lambda}$  is formulated as:

$$Dis^{M,\lambda} = \frac{\sum_{i=1}^N \|\tau_i - \lambda_i \cdot M_i \odot \tau_{uni}\|^2}{N} \leq Dis^M \quad (5)$$

Eq. 5 demonstrates that the distance between the merged model and each individual model can be minimized after applying the rescalers. Please check Appendix B for detailed proof.

### 3.4 Empirical analysis

In Fig. 3, we visualize partial results of merging eight ViT-B/32 models on different tasks using t-SNE [69] and Grad-CAM [61]. It can be seen that each procedure of EMR-MERGING can help improve the performance of the merged model and perform closer to individual models. Specifically, a more obvious distinction is shown in t-SNE and a more precise target is focused by Grad-CAM. Please check Section 4.1.1 for experimental details and Appendix E for more visualization results.

Table 2: Multi-task performance when merging ViT-B/32 models on eight tasks.

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	Avg Acc
Individual	75.3	77.7	96.1	99.7	97.5	98.7	99.7	79.4	90.5
Traditional MTL	73.9	74.4	93.9	98.2	95.8	98.9	99.5	77.9	88.9
Weight Averaging	65.3	63.4	71.4	71.7	64.2	52.8	87.5	50.1	65.8
Fisher Merging [46]	68.6	69.2	70.7	66.4	72.9	51.1	87.9	59.9	68.3
RegMean [33]	65.3	63.5	75.6	78.6	78.1	67.4	93.7	52.0	71.8
Task Arithmetic [30]	63.8	62.1	72.0	77.6	74.4	65.1	94.0	52.2	70.1
Ties-Merging [84]	64.8	62.9	74.3	78.9	83.1	71.4	97.6	56.2	73.6
AdaMerging [85]	64.5	68.1	79.2	93.8	87.0	91.9	97.5	59.1	80.1
AdaMerging++ [85]	66.6	68.3	82.2	94.2	89.6	89.0	98.3	60.6	81.1
<b>EMR-MERGING (Ours)</b>	<b>75.2</b>	<b>72.8</b>	<b>93.5</b>	<b>99.5</b>	<b>96.9</b>	<b>98.1</b>	<b>99.6</b>	<b>74.4</b>	<b>88.7</b>

Table 3: Multi-task performance when merging ViT-L/14 models on eight tasks.

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	Avg Acc
Individual	82.3	92.4	97.4	100	98.1	99.2	99.7	84.1	94.2
Traditional MTL	80.8	90.6	96.3	96.3	97.6	99.1	99.6	84.4	93.5
Weight Averaging	72.1	81.6	82.6	91.9	78.2	70.7	97.1	62.8	79.6
Fisher Merging [46]	69.2	88.6	87.5	93.5	80.6	74.8	93.3	70.0	82.2
RegMean [33]	73.3	81.8	86.1	97.0	88.0	84.2	98.5	60.8	83.7
Task Arithmetic [30]	74.1	82.1	86.7	93.8	87.9	86.8	98.9	65.6	84.5
Ties-Merging [84]	76.5	85.0	89.3	95.7	90.3	83.3	99.0	68.8	86.0
AdaMerging [85]	79.0	90.3	90.8	96.2	93.4	98.0	99.0	79.9	90.8
AdaMerging++ [85]	79.4	90.3	91.6	97.4	93.4	97.5	99.0	79.2	91.0
<b>EMR-MERGING (Ours)</b>	<b>83.2</b>	<b>90.7</b>	<b>96.8</b>	<b>99.7</b>	<b>97.9</b>	<b>99.1</b>	<b>99.7</b>	<b>82.7</b>	<b>93.7</b>

In Fig. 4, we compare sign conflicts, L2 distance, and cosine similarity between the merged model weights obtained by different merging methods and the task-specific model weights. It can be seen that EMR-MERGING significantly reduces sign conflicts and L2 distance and improves the cosine similarity, indicating that EMR-MERGING approximates each task-specific model weight effectively. The configuration of Fig. 4 can be found in Appendix F.




Figure 4: Comparison of (a) sign conflicts, (b) L2 distance, and (c) cosine similarity of model weights obtained by different methods and task-specific model weights.

## 4 Experiment Validation

**Baseline methods.** We compare the proposed EMR-MERGING with: (1) Individual Models, (2) Traditional MTL, (3) Weight Averaging, (4) Fisher Merging [46], (5) RegMean [33], (6) Task Arithmetic [30], (7) Ties-Merging [84], (8) AdaMerging [85]. For more details about baseline methods, please check Appendix C.

### 4.1 Merging vision models

#### 4.1.1 Merging 8 ViTs.

**Settings.** We follow the setting from Task Arithmetic [30], Ties-Merging [84], and AdaMerging [85]. We employ ViT-B/32 and ViT-L/14, two variants of CLIP [54] models’ visual encoders, as the pre-trained models. The performance of each method is evaluated by eight image classification tasks, including SUN397 [83], Cars [35], RESISC45 [10], EuroSAT [27], SVHN [91], GTSRB [65], MNIST [38], and DTD [11]. All the datasets are evaluated by accuracy.

**Results.** The experimental results of merging ViT-B/32 and ViT-L/14 on eight tasks are shown in Tab. 2 and Tab. 3. We observe that EMR-MERGING shows significant performance improvement




Figure 5: Partial visualization results of different merging methods, (a) t-SNE and (b) Grad-CAM.

Table 4: Task-specific and average performance when merging ViT-B/16 models on 30 tasks.

Task-specific Acc	MNIST	Cifar-10	Vegetables	Food-101	Kvasir-v2	Intel-Images	Cars	EuroSAT	Weather	Cats and Dogs
Individual	99.22	97.88	100.00	87.93	94.31	94.63	85.96	99.04	98.22	99.05
Weight Averaging	27.63	42.91	83.20	68.02	25.27	82.40	7.74	24.37	61.06	91.28
RegMean [33]	90.71	89.65	99.10	76.14	71.00	93.60	16.28	74.13	86.62	98.54
Task Arithmetic [30]	30.81	59.86	91.97	73.06	31.05	89.03	9.34	31.25	74.56	93.61
Ties-Merging [84]	23.21	42.82	92.31	73.22	21.09	89.39	5.30	10.98	72.86	91.88
AdaMerging [85]	81.22	87.54	97.97	75.23	22.76	91.02	0.42	44.60	89.13	96.91
<b>EMR-MERGING (Ours)</b>	<b>98.99</b>	<b>96.69</b>	<b>99.97</b>	<b>85.05</b>	<b>93.67</b>	<b>95.27</b>	<b>72.48</b>	<b>96.24</b>	<b>97.76</b>	<b>99.27</b>
	Dogs	Fashion	Pet	LandScape	Flowers	STL-10	CUB-200-2011	EMNIST	DTD	RESISC45
Individual	85.16	93.26	92.23	86.83	98.19	99.07	84.79	94.67	71.76	98.90
Weight Averaging	47.80	20.46	31.26	73.14	68.97	37.74	37.66	7.73	14.63	13.56
RegMean [33]	42.89	83.42	34.62	83.64	95.26	78.94	49.78	48.67	30.53	34.66
Task Arithmetic [30]	47.65	37.11	33.24	79.59	80.68	39.66	41.86	11.05	14.73	15.50
Ties-Merging [84]	26.03	27.05	12.84	78.27	34.33	6.17	31.28	5.61	3.71	6.79
AdaMerging [85]	53.09	76.76	48.34	81.98	95.69	68.91	48.19	18.02	16.68	24.83
<b>EMR-MERGING (Ours)</b>	<b>81.89</b>	<b>92.41</b>	<b>87.15</b>	<b>86.17</b>	<b>97.66</b>	<b>98.41</b>	<b>74.91</b>	<b>92.03</b>	<b>60.05</b>	<b>93.01</b>
	MangoLeafBD	Beans	Cifar-100	GTSRB	SVHN	SUN397	KenyanFood13	Animal-10N	Garbage	Fruits-360
Individual	100.00	97.73	89.85	95.74	96.22	78.98	85.53	92.52	93.36	99.63
Weight Averaging	68.58	70.98	77.98	15.00	10.88	57.42	33.55	46.00	22.89	5.38
RegMean [33]	98.10	92.58	82.59	56.96	66.13	58.58	57.11	68.74	65.31	19.79
Task Arithmetic [30]	87.02	84.62	80.20	37.01	17.41	55.88	36.32	51.14	25.23	6.15
Ties-Merging [84]	76.58	67.22	78.61	40.74	10.54	52.69	19.90	19.13	3.91	1.50
AdaMerging [85]	99.13	93.38	84.19	59.90	25.70	64.09	48.66	66.55	38.54	7.94
<b>EMR-MERGING (Ours)</b>	<b>100.00</b>	<b>98.48</b>	<b>89.09</b>	<b>95.98</b>	<b>82.33</b>	<b>76.19</b>	<b>74.12</b>	<b>87.70</b>	<b>87.11</b>	<b>96.07</b>
Average Acc	Individual	Weight Averaging	RegMean [33]	Task Arithmetic [30]	Ties-Merging [84]	AdaMerging [85]	<b>EMR-MERGING (Ours)</b>			
Acc	93.02	42.52	68.14	48.89	37.53	60.25	<b>89.54</b>			

compared to existing merging methods, respectively 7.6% and 2.7%. Notably, EMR-MERGING requires no additional training, tuning, or any dataset accessibility while outperforming AdaMerging and Ties-Merging, which require additional training or careful hyper-parameter tuning using datasets. Under this setting, EMR-MERGING performs very close to or even better than traditional MTL, which is normally considered as a reference upper bound for model merging work [85]. For visualized comparison, we provide some visualization results of different merging methods using t-SNE and Grad-CAM in Fig. 5. It can be seen that among all the merging methods, the visualization results of EMR-MERGING are the closest to individual models, which corresponds to quantitative results. Please check Appendix E for more visualization results.

#### 4.1.2 Merging 30 ViTs.

**Settings.** To further explore the performance of EMR-MERGING, we establish a new benchmark on merging vision models, expanding the number of task-specific models from eight to 30. We employ ViT-B/16 [22] pre-trained on ImageNet-21k [18] as the pre-trained model. The performance is evaluated by image classification datasets including MNIST [38], CIFAR-10 [36], Vegetables [1], Food-101 [6], Kvasir-v2 [53], Cars [35], Intel Images [4], EuroSAT [27], Weather [82], Cats and dogs [15], MangoLeafBD [2], Beans [37], CIFAR-100 [36], GTSRB [65], SVHN [91], Dogs [34], Fashion MNIST [81], Oxford-IIIT-Pet [50], Landscape Recognition [17], Flowers Recognition [45], STL-10 [12], CUB-200-2011 [73], EMNIST [13], DTD [11], RESISC45 [10], SUN397 [83], Kenyan-Food13 [32], Animal-10N [64], Garbage Classification [8], and Fruits-360 [47], covering tasks from common food classification to disease detection. All of them are evaluated by accuracy.

**Results.** The experimental results are shown in Tab. 4. It can be clearly seen that under this challenging setting of merging 30 models, all the existing methods show significant performance drops compared to individual models. Even RegMean, which performs best among existing methods, still exhibits a performance decay of nearly 25%. However, EMR-MERGING can reduce this value

Table 5: Results of merging RoBERTa models on eight datasets from GLUE benchmark.

Methods	Single-Sentence Tasks		Similarity and Paraphrase Tasks			Inference Tasks		
	CoLA	SST2	MRPC	STSB	QQP	MNLI	QNLI	RTE
Individual	0.6018	0.9404	0.8922	0.9063	0.9141	0.8720	0.9271	0.7906
Weight Averaging	0.1396	0.6411	0.6936	0.3184	0.7536	0.4219	0.587	0.5523
RegMean [33]	0.3667	0.906	0.7574	0.6268	0.8355	0.7002	0.8235	0.5848
Task Arithmetic [30]	0.1878	0.8589	0.7990	0.7403	0.8378	0.5908	0.6967	0.6209
Ties-Merging [84]	0.2048	0.8440	0.8113	0.5819	0.8570	0.6465	0.7481	0.4296
<b>EMR-MERGING (Ours)</b>	<b>0.3996</b>	<b>0.9335</b>	<b>0.8627</b>	<b>0.8277</b>	<b>0.8972</b>	<b>0.8545</b>	<b>0.8957</b>	<b>0.7437</b>

Table 6: Multi-task performance when merging GPT-2 models on seven text classification tasks.

Method	CoLA	MNLI	MRPC	QNLI	QQP	RTE	SST-2	Avg.
Individual	76.8	82.1	80.4	88.3	89.6	65.3	91.2	82.0
Weight Averaging	55.0	55.1	51.0	57.6	76.7	44.8	52.5	56.1
Fisher Merging [46]	54.8	58.0	39.5	63.3	81.5	49.1	64.7	58.7
RegMean [33]	61.7	70.4	65.4	69.7	78.8	56.0	79.7	68.8
Task Arithmetic [30]	68.7	68.6	69.6	70.5	81.8	47.3	83.6	70.0
Ties-Merging [84]	68.4	71.4	68.4	69.6	82.4	47.7	81.8	70.0
<b>EMR-MERGING (Ours)</b>	<b>72.8</b>	<b>81.1</b>	<b>79.2</b>	<b>84.8</b>	<b>88.1</b>	<b>66.5</b>	<b>90.3</b>	<b>80.4</b>

to 3.48%. This shows that the proposed method maintains the performance comparable to individual models when merging vision models even if the number of tasks increases.

## 4.2 Merging language models

### 4.2.1 Merging fully finetuned RoBERTa models

**Settings.** We partially follow the setting from DARE [90]. However, instead of merging two or three models at a time, we merge all eight models finetuned on each task. RoBERTa-base [43] model is selected as the pre-trained model. The performance of each method is evaluated by eight tasks from GLUE [74] benchmark, respectively CoLA [76], SST-2 [63], MRPC [21], STS-B [9], QQP [31], MNLI [78], QNLI [56], and RTE [24]. Among them, CoLA is evaluated by the Matthews correlation coefficient, STS-B is evaluated by the average value of Pearson and Spearman correlation coefficients, and the rest tasks are evaluated by accuracy.

**Results.** The experimental results are shown in Tab. 5. It can be seen that EMR-MERGING outperforms all the existing methods on every task, verifying the applicability of the proposed method to language models. Note that the reported results of Ties-Merging, Task Arithmetic, and RegMean are the best among multiple hyper-parameter settings. Please check Appendix D.4 for more detailed information. It should also be noted that we find that under our setting of merging multiple models, DARE may not help improve the performance. Similar results were also reported by [25]. This may be due to DARE’s random dropping strategy can no longer resolve conflicts among task vectors under the setting of merging multiple models. Please check Appendix D.3 for DARE’s experimental results.

### 4.2.2 Merging fully finetuned GPT-2 models

**Settings.** We follow the setting from FusionBench [68], a benchmark for model merging. We merge GPT-2 [55] models on seven tasks from GLUE [74], each with a different head for classification. Under this setting, each task is evaluated by accuracy.

**Results.** The experimental results are shown in Tab. 6. EMR-MERGING outperforms all the merging methods by over 10% and decreases the performance degradation caused by model merging from 12% to 1.6%. This validates the applicability of EMR-MERGING to fully finetuned GPT2-scale language models.

### 4.2.3 Merging PEFT models

**Settings.** We follow the setting from Ties-Merging [84]. (IA)<sup>3</sup> [42] is a PEFT method that uses learned vectors to scale the base model activations. We use T0-3B [60] as the base model and merge (IA)<sup>3</sup> modules. The performance is evaluated using eleven datasets, including RTE [24], CB [16],



Table 7: Results of merging (IA)<sup>3</sup> models on eleven NLP tasks.

Methods	Validation	RTE	CB	Winogrande	WiC	WSC	COPA	H-SWAG	Story Cloze	ANLI-R1	ANLI-R2	ANLI-R3	Avg Acc
Individual	-	82.7	95.8	75.1	71.7	65.3	85.3	44.4	94.9	70.2	46.5	53	71.4
Traditional MTL	-	88.6	95.8	75.5	61.1	80.6	94.1	42.3	97.6	70.5	49.8	47.7	73.1
Fisher Merging [46]	✓	83.3	83.3	56.7	54.2	58.3	83.1	42.2	94.1	45.9	41.0	42.2	62.2
RegMean [33]	✓	81.2	58.3	53.8	55.2	53.5	80.9	40.1	92.5	43.3	39.2	40.2	58
Task Arithmetic [30]	✓	74.1	83.3	62.8	49.1	49.3	87.5	41.5	95.3	60.8	49.4	50.0	63.9
Ties-Merging [84]	✓	78.0	83.3	67.9	57.6	59.7	81.7	42.8	90.3	66.9	51.3	51.1	66.4
Weight Averaging	×	81.2	58.3	53.8	55.2	53.5	80.9	40.1	92.5	43.3	39.2	40.2	58
Task Arithmetic [30]	×	76.5	79.2	57.7	51.6	51.4	66.2	31.4	81.5	59.8	47.5	48.2	59.2
Ties-Merging [84]	×	81.2	87.5	60.8	59.9	58.3	80.2	42.6	91.1	58.1	46.5	47.4	64.9
<b>EMR-MERGING (Ours)</b>	×	<b>81.8</b>	<b>87.5</b>	<b>66.6</b>	<b>56.1</b>	<b>65.3</b>	<b>82.4</b>	<b>44.7</b>	<b>93.6</b>	<b>65.7</b>	<b>43.8</b>	<b>50.8</b>	<b>67.1</b>

Table 8: Results of merging multi-modal BEiT3 models on five vision-language tasks.

Methods	Task Metric	COCO-Captioning					ImageNet-1k Classification	NLVR2	VQAv2
		COCO-Retrieval Accuracy(↑)	BLEU4(↑)	CIDEr(↑)	METEOR(↑)	ROUGE-L(↑)	Accuracy(↑)	Accuracy(↑)	Accuracy(↑)
Individual		0.8456	0.394	1.337	0.311	0.601	0.8537	0.7765	0.8439
Weight Averaging		0.1893	0.031	0.001	0.115	0.159	0.6771	0.2800	0.6285
Task Arithmetic [30]		0.3177	0.033	0.000	0.118	0.176	0.7081	0.3809	0.6933
Ties-Merging [84]		0.3929	0.029	0.001	0.108	0.167	0.6978	0.3206	0.6717
<b>EMR-MERGING(Ours)</b>		<b>0.7946</b>	<b>0.289</b>	<b>1.060</b>	<b>0.272</b>	<b>0.534</b>	<b>0.7742</b>	<b>0.7475</b>	<b>0.7211</b>

Winogrande [59], WiC [52], WSC [39], COPA [58], H-SWAG [92], Story Cloze [62], and ANLI [48] from R1 to R3. All the datasets are evaluated by accuracy.

**Results.** The experimental results are shown in Tab. 7. EMR-MERGING outperforms all the merging methods. Compared to methods without validation, EMR-MERGING improves the average accuracy on each task by 2.2%. When compared to methods that require validation data to tune hyper-parameters or compute matrices for weighted merging, EMR-MERGING still improves the average performance by 0.7%, validating the applicability of our method to PEFT models.

### 4.3 Merging multi-modal models

**Settings.** We merge BEiT3-base [75] models finetuned on five datasets from different kinds of tasks, respectively ImageNet-1k [18] (Image Classification), VQAv2 [26] (Visual Question Answering), NLVR2 [67] (Visual Reasoning), COCO Captioning [41] (Image Captioning), and COCO Retrieval [41] (Image-Text Retrieval). Among them, COCO Captioning is evaluated by BLEU4 [49], CIDEr [72], METEOR [3], and ROUGE-L [40]. The other tasks are evaluated by accuracy.

**Results.** The experimental results are shown in Tab. 8. It can be seen that EMR-MERGING performs best on all the vision-language tasks regardless of which evaluation metric is applied among all the merging methods, validating the effectiveness of EMR-MERGING in merging multi-modal models.

### 4.4 Merging different number of models

In Fig. 6, we compare the number of parameters and performance using individual models, Ties-Merging, and EMR-MERGING when merging different numbers of ViT-B/32 models under the setting of no-validation.

**Number of parameters.** Compared to other merging methods, EMR-MERGING requires a little additional storage for task-specific modulators. However, compared to a single 32-bit model, the additional storage caused by a task-specific 1-bit mask equals a binarized network, whose size is 32 times smaller than a single 32-bit model [14]. Additionally, the storage required by a task-specific rescaler, which is a single parameter, is negligible. In Fig. 6(a), we compare the number of parameters when merging different numbers of models, and we observe that EMR-MERGING’s parameter number is slightly more than Ties-Merging but significantly fewer than individual models.




Figure 6: Comparison of the (a) number of parameters and (b) average normalized performance when using individual models, Ties-Merging, and EMR-MERGING.

Table 9: Ablation on the Electing procedure of EMR-MERGING.

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	Avg Acc
Task Arithmetic	63.8	62.1	72.0	77.6	74.4	65.1	94.0	52.2	70.1
Task Arithmetic w/ M&R	<b>67.6</b>	<b>67.3</b>	<b>80.2</b>	<b>91.3</b>	<b>79.3</b>	<b>75.7</b>	<b>96.0</b>	<b>57.9</b>	<b>76.9</b> [↑ 6.8]
Ties-Merging	64.8	62.9	74.3	78.9	83.1	71.4	97.6	56.2	73.6
Ties-Merging w/ M&R	<b>68.8</b>	<b>68.9</b>	<b>82.2</b>	<b>91.6</b>	<b>81.4</b>	<b>80.0</b>	<b>96.6</b>	<b>59.3</b>	<b>78.6</b> [↑ 5.0]
AdaMerging++	66.6	68.3	82.2	94.2	89.6	89.0	98.3	60.6	81.1
AdaMerging++ w/ M&R	<b>74.0</b>	<b>76.2</b>	<b>93.1</b>	<b>98.2</b>	<b>93.3</b>	<b>96.3</b>	<b>99.4</b>	<b>71.2</b>	<b>87.7</b> [↑ 6.6]
EMR-MERGING (Ours)	<b>75.2</b>	<b>72.8</b>	<b>93.5</b>	<b>99.5</b>	<b>96.9</b>	<b>98.1</b>	<b>99.6</b>	<b>74.4</b>	<b>88.7</b>

Table 10: Ablation on the Masking and Rescaling procedures of EMR-MERGING.

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	Avg Acc
Ours (Elect)	31.7	34.7	51.8	65.9	85.7	64.0	98.2	42.2	59.3
Ours (Elect & Mask)	70.7	65.9	92.2	98.7	96.9	97.6	99.6	72.3	86.8 [↑ 27.5]
Ours (Elect & Rescale)	58.2	57.2	69.1	81.6	85.2	73.0	98.4	52.2	71.9 [↑ 12.6]
Ours (Elect, Mask & Rescale)	<b>75.2</b>	<b>72.8</b>	<b>93.5</b>	<b>99.5</b>	<b>96.9</b>	<b>98.1</b>	<b>99.6</b>	<b>74.4</b>	<b>88.7</b> [↑ 29.4]

**Performance.** The performance comparison when merging different numbers of models is shown in Fig. 6(b). Compared to Ties-Merging, the performance of EMR-MERGING is higher and decreases more slowly as the task increases. Note that EMR-MERGING outperforms individual models under the 2-task setting. Similar findings are reported by DARE [90]. More details are shown in Appendix D.5.

#### 4.5 Ablation Study

We perform ablations on all the components of EMR-MERGING as follows.

**Ablation on Electing procedure.** Tab. 9 shows the results of merging eight ViT-B/32 models when the Electing procedure is replaced by other task vector-based merging methods. The effectiveness of our Electing strategy is verified by outperforming the combination of other merging methods with masking and rescaling. Another interesting finding is that as a post-processing procedure, masking and rescaling can help improve the performance of task vector-based merging methods, respectively 6.8%, 5.0%, and 6.6% for Task Arithmetic, Ties-Merging, and AdaMerging++.

**Ablation on Masking and Rescaling procedures.** Then, we further validate the importance of Masking and Rescaling procedures by disabling either or both of them. The results are shown in Tab. 10. It can be seen that simply electing results in a severe performance drop while adding Masking and Rescaling can improve the performance by 27.5% and 12.6%, respectively. Furthermore, compared to separately applying either of these two procedures, jointly applying Masking and Rescaling leads to greater improvement, up to 29.4%.

## 5 Conclusion

In this paper, we study on tuning-free and high-performance model merging. We first attribute the severe performance degradation of existing merging methods to that a single model can hardly simulate all the models’ performance. Then we propose ELECT, MASK & RESCALE-MERGING (EMR-MERGING), which does not require any data access or additional training for tuning. The effectiveness of EMR-MERGING is validated by comprehensive experiments on various classical benchmarks and newly-established benchmarks under vision, NLP, PEFT, and multi-modal settings.

## 6 Acknowledgement

This work is supported by Shanghai Natural Science Foundation (No. 23ZR1402900), National Natural Science Foundation of China (No. 62071127 and No. 62306261), National Key Research and Development Program of China (No. 2022ZD0160101). The computations in this research were performed using the CFFF platform of Fudan University.

## References

- [1] M. I. Ahmed, S. M. Mamun, and A. U. Z. Asif. Dcnn-based vegetable image classification using transfer learning: A comparative study. In *2021 5th International Conference on Computer, Communication and Signal Processing (ICCCSP)*, pages 235–243. IEEE, 2021.
- [2] S. I. Ahmed, M. Ibrahim, M. Nadim, M. M. Rahman, M. M. Shejunti, T. Jabid, and M. S. Ali. Mangoleafbd: A comprehensive image dataset to classify diseased and healthy mango leaves. *Data in Brief*, 47:108941, 2023.
- [3] S. Banerjee and A. Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.
- [4] P. Bansal. Intel image classification. Available on <https://www.kaggle.com/puneet6060/intel-image-classification>, Online, 2019.
- [5] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [6] L. Bossard, M. Guillaumin, and L. Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014.
- [7] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [8] CCHANG. Garbage classification. <https://www.kaggle.com/ds/81794>, 2018.
- [9] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, and L. Specia. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*, 2017.
- [10] G. Cheng, J. Han, and X. Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017.
- [11] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, and A. Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014.
- [12] A. Coates, A. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 215–223. JMLR Workshop and Conference Proceedings, 2011.
- [13] G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE, 2017.
- [14] M. Courbariaux, I. Hubara, D. Soudry, R. El-Yaniv, and Y. Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- [15] W. Cukierski. Dogs vs. cats, 2013. URL <https://kaggle.com/competitions/dogs-vs-cats>.
- [16] M.-C. De Marneffe, M. Simons, and J. Tonhauser. The commitmentbank: Investigating projection in naturally occurring discourse. In *proceedings of Sinn und Bedeutung*, volume 23, pages 107–124, 2019.
- [17] DeepNets. Landscape recognition. <https://www.kaggle.com/datasets/utkarshsaxenadn/landscape-recognition-image-dataset-12k-images>.
- [18] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [19] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [20] J. Dodge, G. Ilharco, R. Schwartz, A. Farhadi, H. Hajishirzi, and N. Smith. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping. *arXiv preprint arXiv:2002.06305*, 2020.
- [21] B. Dolan and C. Brockett. Automatically constructing a corpus of sentential paraphrases. In *Third international workshop on paraphrasing (IWP2005)*, 2005.

- [22] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [23] R. A. Fisher. On the mathematical foundations of theoretical statistics. *Philosophical transactions of the Royal Society of London. Series A, containing papers of a mathematical or physical character*, 222(594-604):309–368, 1922.
- [24] D. Giampiccolo, B. Magnini, I. Dagan, and W. B. Dolan. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pages 1–9, 2007.
- [25] C. Goddard, S. Siriwardhana, M. Ehghaghi, L. Meyers, V. Karpukhin, B. Benedict, M. McQuade, and J. Solawetz. Arcee’s mergekit: A toolkit for merging large language models. *arXiv preprint arXiv:2403.13257*, 2024.
- [26] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, and D. Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- [27] P. Helber, B. Bischke, A. Dengel, and D. Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019.
- [28] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019.
- [29] E. J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021.
- [30] G. Ilharco, M. T. Ribeiro, M. Wortsman, L. Schmidt, H. Hajishirzi, and A. Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*, 2022.
- [31] S. Iyer, N. Dandekar, K. Csernai, et al. First quora dataset release: Question pairs. data. quora. com. 2017.
- [32] M. Jalal, K. Wang, S. Jefferson, Y. Zheng, E. O. Nsoesie, and M. Betke. Scraping social media photos posted in kenya and elsewhere to detect and analyze food types. In *Proceedings of the 5th International Workshop on Multimedia Assisted Dietary Management*, pages 50–59, 2019.
- [33] X. Jin, X. Ren, D. Preotiuc-Pietro, and P. Cheng. Dataless knowledge fusion by merging weights of language models. In *The Eleventh International Conference on Learning Representations*, 2022.
- [34] A. Khosla, N. Jayadevaprakash, B. Yao, and L. Fei-Fei. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, June 2011.
- [35] J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013.
- [36] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [37] M. A. Lab. Bean disease dataset, January 2020.
- [38] Y. LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [39] H. Levesque, E. Davis, and L. Morgenstern. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*, 2012.
- [40] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [41] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [42] H. Liu, D. Tam, M. Muqeeth, J. Mohta, T. Huang, M. Bansal, and C. A. Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.

- [43] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [44] T. maintainers and contributors. Torchvision: Pytorch’s computer vision library. <https://github.com/pytorch/vision>, 2016.
- [45] A. Mamaev. Flowers recognition. <https://www.kaggle.com/datasets/alxmamaev/flowers-recognition>.
- [46] M. S. Matena and C. A. Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022.
- [47] H. Muresan and M. Oltean. Fruit recognition from images using deep learning. *Acta Universitatis Sapientiae, Informatica*, 10(1):26–42, 2018.
- [48] Y. Nie, A. Williams, E. Dinan, M. Bansal, J. Weston, and D. Kiela. Adversarial nli: A new benchmark for natural language understanding. *arXiv preprint arXiv:1910.14599*, 2019.
- [49] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
- [50] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012.
- [51] S. Paul and P.-Y. Chen. Vision transformers are robust learners. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 36, pages 2071–2081, 2022.
- [52] M. T. Pilehvar and J. Camacho-Collados. Wic: the word-in-context dataset for evaluating context-sensitive meaning representations. *arXiv preprint arXiv:1808.09121*, 2018.
- [53] K. Pogorelov, K. R. Randel, C. Griwodz, S. L. Eskeland, T. de Lange, D. Johansen, C. Spampinato, D.-T. Dang-Nguyen, M. Lux, P. T. Schmidt, et al. Kvasir: A multi-class image dataset for computer aided gastrointestinal disease detection. In *Proceedings of the 8th ACM on Multimedia Systems Conference*, pages 164–169, 2017.
- [54] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [55] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.
- [56] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- [57] A. Ramé, K. Ahuja, J. Zhang, M. Cord, L. Bottou, and D. Lopez-Paz. Model ratatouille: Recycling diverse models for out-of-distribution generalization. In *International Conference on Machine Learning*, pages 28656–28679. PMLR, 2023.
- [58] M. Roemmele, C. A. Bejan, and A. S. Gordon. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *2011 AAAI Spring Symposium Series*, 2011.
- [59] K. Sakaguchi, R. L. Bras, C. Bhagavatula, and Y. Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- [60] V. Sanh, A. Webson, C. Raffel, S. H. Bach, L. Sutawika, Z. Alyafeai, A. Chaffin, A. Stiegler, T. L. Scao, A. Raja, et al. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*, 2021.
- [61] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [62] R. Sharma, J. Allen, O. Bakhshandeh, and N. Mostafazadeh. Tackling the story ending biases in the story cloze test. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 752–757, 2018.

- [63] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [64] H. Song, M. Kim, and J.-G. Lee. Selfie: Refurbishing unclean samples for robust deep learning. In *International conference on machine learning*, pages 5907–5915. PMLR, 2019.
- [65] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel. The german traffic sign recognition benchmark: a multi-class classification competition. In *The 2011 international joint conference on neural networks*, pages 1453–1460. IEEE, 2011.
- [66] G. Stoica, D. Bolya, J. Bjorner, T. Hearn, and J. Hoffman. Zipit! merging models from different tasks without training. *arXiv preprint arXiv:2305.03053*, 2023.
- [67] A. Suhr, S. Zhou, A. Zhang, I. Zhang, H. Bai, and Y. Artzi. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*, 2018.
- [68] A. Tang, L. Shen, Y. Luo, H. Hu, B. Du, and D. Tao. FusionBench: A Comprehensive Benchmark of Deep Model Fusion, June 2024.
- [69] L. van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.
- [70] S. Vandenhende, S. Georgoulis, W. Van Gansbeke, M. Proesmans, D. Dai, and L. Van Gool. Multi-task learning for dense prediction tasks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3614–3633, 2021.
- [71] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [72] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015.
- [73] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [74] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [75] W. Wang, H. Bao, L. Dong, J. Bjorck, Z. Peng, Q. Liu, K. Aggarwal, O. K. Mohammed, S. Singhal, S. Som, and F. Wei. Image as a foreign language: BEiT pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [76] A. Warstadt, A. Singh, and S. R. Bowman. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641, 2019.
- [77] R. Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- [78] A. Williams, N. Nangia, and S. R. Bowman. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*, 2017.
- [79] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, et al. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.
- [80] M. Wortsman, G. Ilharco, S. Y. Gadre, R. Roelofs, R. Gontijo-Lopes, A. S. Morcos, H. Namkoong, A. Farhadi, Y. Carmon, S. Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022.
- [81] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [82] H. Xiao, F. Zhang, Z. Shen, K. Wu, and J. Zhang. Classification of weather phenomenon from images by using deep convolutional neural network. *Earth and Space Science*, 8(5):e2020EA001604, 2021.

- [83] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3485–3492. IEEE, 2010.
- [84] P. Yadav, D. Tam, L. Choshen, C. Raffel, and M. Bansal. Ties-merging: Resolving interference when merging models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [85] E. Yang, Z. Wang, L. Shen, S. Liu, G. Guo, X. Wang, and D. Tao. Adamerging: Adaptive model merging for multi-task learning. In *The Twelfth International Conference on Learning Representations*, 2023.
- [86] P. Ye, T. He, S. Tang, B. Li, T. Chen, L. Bai, and W. Ouyang. Stimulative training++: Go beyond the performance limits of residual networks. *arXiv preprint arXiv:2305.02507*, 2023.
- [87] P. Ye, C. Huang, M. Shen, T. Chen, Y. Huang, Y. Zhang, and W. Ouyang. Merging vision transformers from different tasks and domains. *arXiv preprint arXiv:2312.16240*, 2023.
- [88] P. Ye, B. Li, Y. Li, T. Chen, J. Fan, and W. Ouyang. b-darts: Beta-decay regularization for differentiable architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10874–10883, June 2022.
- [89] P. Ye, S. Tang, B. Li, T. Chen, and W. Ouyang. Stimulative training of residual networks: A social psychology perspective of loafing. *Advances in Neural Information Processing Systems*, 35:3596–3608, 2022.
- [90] L. Yu, B. Yu, H. Yu, F. Huang, and Y. Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. *arXiv preprint arXiv:2311.03099*, 2023.
- [91] N. Yuval. Reading digits in natural images with unsupervised feature learning. In *Proceedings of the NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [92] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- [93] B. Zhang, J. Yuan, B. Shi, T. Chen, Y. Li, and Y. Qiao. Uni3d: A unified baseline for multi-dataset 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9253–9262, 2023.
- [94] J. Zhang, J. Liu, J. He, et al. Composing parameter-efficient modules with arithmetic operation. *Advances in Neural Information Processing Systems*, 36:12589–12610, 2023.
- [95] Y. Zhang and Q. Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609, 2021.

# Appendix for EMR-MERGING

## A Algorithm flow of EMR-MERGING

We summarize the procedure of EMR-MERGING in Algorithm 1.

---

### Algorithm 1 EMR-MERGING Procedure

---

**Input:** Finetuned models  $W_{1..N}$ , pretrained model  $W_{pre}$

**Output:** Unified task vector  $\tau_{uni}$ , task-specific masks  $M_{1..N}$ , task-specific rescalers  $\lambda_{1..N}$

```

for  $t$  in  $1, \dots, N$  do
  ▷ Create task vectors.
   $\tau_t = W_t - W_{pre}$ 
end
▷ Step 1: Elect the unified task vector.
 $\gamma_{uni} = \text{sgn}(\sum_{t=1}^n \tau_t)$ 
 $\epsilon_{uni} = \text{zeros}(d)$ 
for  $t$  in  $1, \dots, N$  do
  for  $p$  in  $1, \dots, d$  do
    if  $\gamma_{uni}^p \cdot \tau_t^p > 0$  then
       $\epsilon_{uni}^p = \max(\epsilon_{uni}^p, \text{abs}(\gamma_{uni}^p))$ 
    end
  end
end
 $\tau_{uni} = \gamma_{uni} \odot \epsilon_{uni}$ 
for  $t$  in  $1, \dots, N$  do
  ▷ Step 2: Generate task-specific masks.
  for  $p$  in  $1, \dots, d$  do
     $M_t^p = \text{bool}(\tau_t^p \odot \tau_{uni}^p > 0)$ 
  end
  ▷ Step 3: Generate task-specific rescalers.
   $\lambda_t = \frac{\text{sum}(\text{abs}(\tau_t))}{\text{sum}(\text{abs}(M_t \cdot \tau_{uni}))}$ 
end

```

---

## B Theoretical analyses

In Section 3, we claimed that the task-specific modulators can lower the distance between the merged model and task-specific models. Here we provide detailed theoretical analyses.

Our goal is to merge model weights  $W_{1..N}$  by minimizing the distance between the merged model  $W_{uni}$  and each individual models  $W_i$ ,  $i \in [1..N]$  **without** using any dataset  $[X_i, Y_i]$ , where the distance can be calculated by:

$$Dis = \frac{\sum_{i=1}^N \|W_i - W_{uni}\|^2}{N} \quad (6)$$

The premise of merging is that all the models are fine-tuned from the same pre-trained model. Thus, Eq. 6 can be re-written:

$$Dis = \frac{\sum_{i=1}^N \|\tau_i - \tau_{uni}\|^2}{N} \quad (7)$$

where  $\tau_i$  refers to the task vector for Task  $i$ .  $\tau_{uni}$  is the merged task vector. We demonstrate the effectiveness of the task-specific modulators by step.

**Analysis 1: Effectiveness of Masks.** Suppose we apply a mask  $M_i$  to the unified model  $\tau_{uni}$  to disable elements in  $\tau_{uni}$  that have the opposite sign of the corresponding elements in  $\tau_{uni}$ , which can be written as:






Figure 7: Comparison of (a) sign conflicts, (b) L2 distance, and (c) cosine similarity of model weights obtained by different methods (including AdaMerging++ and each procedure of EMR-MERGING) and task-specific model weights. The detailed configuration is shown in Appendix F.

Table 11: Multi-task performance when merging ViT-B/16 models on eight tasks.

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	Avg Acc
Task Arithmetic [30]	61.1	65.9	74.0	76.2	88.0	73.9	98.4	53.0	73.8
Ties-Merging [84]	69.1	72.5	80.5	84.0	85.0	71.5	98.1	54.9	77.0
AdaMerging [85]	70.2	80.7	81.6	94.8	91.6	95.8	98.5	66.2	84.9
AdaMerging++ [85]	71.8	80.8	84.1	94.3	91.9	94.5	98.7	69.8	85.7
<b>EMR-MERGING (Ours)</b>	<b>78.6</b>	<b>82.6</b>	<b>95.5</b>	<b>99.2</b>	<b>97.6</b>	<b>98.8</b>	<b>99.6</b>	<b>78.3</b>	<b>91.3</b>

$$M_i = (\tau_i \odot \tau_{uni} > 0) \quad (8)$$

By applying the masks  $M_i, i \in [1..N]$ , the distance becomes:

$$Dis^M = \frac{\sum_{i=1}^N \|\tau_i - M_i \odot \tau_{uni}\|^2}{N} \quad (9)$$

Furthermore, it can be written as:

$$\begin{aligned} Dis^M &= \frac{\sum_{i=1}^N \|M_i \odot \tau_i - M_i \odot \tau_{uni}\|^2}{N} + \frac{\sum_{i=1}^N \|(1 - M_i) \odot \tau_i\|^2}{N} \\ &= \frac{\sum_{i=1}^N \|M_i \odot (abs(\tau_i) - abs(\tau_{uni}))\|^2}{N} + \frac{\sum_{i=1}^N \|(1 - M_i) \odot abs(\tau_i)\|^2}{N} \end{aligned} \quad (10)$$

where  $abs(\cdot)$  returns the absolute value of each element in the input. For ease of comparison, the distance without applying  $M_i$  can be formulated as:

$$\begin{aligned} Dis &= \frac{\sum_{i=1}^N \|M_i \odot (abs(\tau_i) - abs(\tau_{uni}))\|^2}{N} + \frac{\sum_{i=1}^N \|(1 - M_i) \odot (abs(\tau_i) + abs(\tau_{uni}))\|^2}{N} \\ &= Dis^M + \frac{\sum_{i=1}^N \|(1 - M_i) \odot abs(\tau_{uni})\|^2}{N} \end{aligned} \quad (11)$$

Thus, we demonstrate that  $Dis^M \leq Dis$ , indicating applying task-specific masks can reduce the distance between the merged model and individual models, thus showing effectiveness.

**Analysis 2: Effectiveness of Rescalers.** Suppose we apply a rescaler  $\lambda_i > 0$  to the masked unified task vector  $M_i \cdot \tau_{uni}$ , the distance becomes:

$$\begin{aligned} Dis^{M,\lambda} &= \frac{\sum_{i=1}^N \|\tau_i - \lambda_i \cdot M_i \odot \tau_{uni}\|^2}{N} \\ &= \frac{\sum_{i=1}^N \|abs(\tau_i) - \lambda_i \cdot abs(M_i \odot \tau_{uni})\|^2}{N} \end{aligned} \quad (12)$$



Figure 8: t-SNE visualization results of different merging methods.

Table 12: Multi-task performance when merging ViT-B/32 models on 9 vision tasks (ImageNet-1K added).

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	ImageNet-1K	Avg Acc
Individual	75.3	77.7	96.1	99.7	97.5	98.7	99.7	79.4	82.0	89.6
Weight Averaging	61.8	56.4	65.9	66.2	62.7	44.5	81.8	49.0	61.5	61.1
Task Arithmetic [30]	51.8	30.9	55.8	64.3	69.0	42.2	92.7	46.8	66.6	57.8
Ties-Merging [84]	53.3	34.1	57.0	55.8	72.3	43.2	90.5	46.5	68.9	58.0
<b>EMR-MERGING (Ours)</b>	<b>77.0</b>	<b>75.2</b>	<b>92.9</b>	<b>92.7</b>	<b>79.7</b>	<b>90.2</b>	<b>97.6</b>	<b>76.2</b>	<b>79.8</b>	<b>84.6</b>

To minimize the distance in Eq. 12, we set the first derivative of  $Dis^\lambda$  with respect to  $\lambda_i$  to 0, thus  $\lambda_i$  can be calculated by:

$$\lambda_i = \frac{\text{sum}(\text{abs}(\tau_i))}{\text{sum}(\text{abs}(M_i \odot \tau_{uni}))} \quad (13)$$

which exactly matches our setting of  $\lambda_i$ . This indicates that our setting of rescalers  $\lambda_i$  can minimize the distance between the merged model and individual models, which is:  $Dis^{M,\lambda} \leq Dis^M$ , thus showing effectiveness.

It is also reflected in Fig. 7 that after Masking and Rescaling, the sign conflicts and L2 distance between the merged model and task-specific models are reduced and the cosine similarity can be improved.




Figure 9: Grad-CAM visualization results of different merging methods.

## C Baseline Methods

- **Individual Models** refer to task-specific models before merging.
- **Traditional MTL** uses datasets from all the tasks to train a single model jointly.
- **Weight Averaging** element-wisely averages all the model weights. Its effectiveness when applied to fine-tuned model weights from the same pre-training has been verified [80, 57, 33].
- **Fisher Merging** [46] uses Fisher information matrices [23] to calculate the importance of each parameter and weighted merges them based on their importance.
- **RegMean** [33] weighted merges models based on a closed-form solution to the merging problem. When merging  $K$  linear model weights  $W_i$ , where  $f_i(x) = W_i^T x$ ,  $i = 1..K$ , the merging problem can be formulated as:  $\min_W \sum_{i=1}^K \|W^T X_i - W_i^T X_i\|^2$ , where  $W$  is the merged model weights, and  $X_i$  denotes the input of  $i^{th}$  model. The closed-form solution to the problem is:  $W = (\sum_{i=1}^K X_i^T X_i)^{-1} (\sum_{i=1}^K X_i^T X_i W_i)$ . Inner-product matrices need to be computed before merging.
- **Task Arithmetic** [30] defines task vectors as the difference between finetuned model weights and the pre-trained model weights. Suppose a model  $\theta_i$  is finetuned from a pre-trained model  $\theta_{pre}$ , the task vector is  $\tau_i = \theta_i - \theta_{pre}$ . When merging  $\theta_{1..K}$ , the merged model is  $\theta_M = \lambda \sum_{i=1}^K \tau_i + \theta_{pre}$ , where  $\lambda$  is the merging coefficient.
- **Ties-Merging** [84] (Trim, Elect Sign & Merge) believes that the conflicts among the task vectors severely effect the merged model's performance. Ties-Merging solves this problem by eliminating redundant parameters and resolving symbol conflicts.

Table 13: Performance of RegMean and Task Arithmetic when pre-processed using DARE [90].

Methods	Single-Sentence Tasks		Similarity and Paraphrase Tasks			Inference Tasks		
	CoLA	SST2	MRPC	STSB	QQP	MNLI	QNLI	RTE
Individual	0.6018	0.9404	0.8922	0.9063	0.9141	0.8720	0.9271	0.7906
<b>EMR-MERGING (Ours)</b>	<b>0.3996</b>	<b>0.9335</b>	<b>0.8627</b>	<b>0.8277</b>	<b>0.8972</b>	<b>0.8545</b>	<b>0.8957</b>	<b>0.7437</b>
RegMean [33]	0.3667	0.906	0.7574	0.6268	0.8355	0.7002	0.8235	0.5848
w/ DARE (drop 10%)	0.5046	0.5298	0.3603	0.1533	0.4955	0.3245	0.4924	0.4477
w/ DARE (drop 30%)	0.4535	0.6135	0.3186	0.0471	0.4219	0.3325	0.505	0.5126
w/ DARE (drop 50%)	0.2758	0.5138	0.3211	-0.0965	0.3685	0.3338	0.508	0.5235
w/ DARE (drop 70%)	0	0.4908	0.3162	0.0021	0.3682	0.3184	0.5056	0.4838
w/ DARE (drop 90%)	0	0.4908	0.3162	-0.0776	0.3682	0.3187	0.5158	0.4910
Task Arithmetic [30]	0.1878	0.8589	0.7990	0.7403	0.8378	0.5908	0.6967	0.6209
w/ DARE (drop 10%)	0.2424	0.8509	0.7966	0.7234	0.8382	0.5869	0.7368	0.6101
w/ DARE (drop 30%)	0.3040	0.8452	0.7941	0.6311	0.8333	0.5515	0.786	0.6137
w/ DARE (drop 50%)	0.2451	0.8188	0.7990	0.4262	0.8099	0.4591	0.7269	0.6029
w/ DARE (drop 70%)	0	0.7225	0.6373	0.1353	0.7321	0.3453	0.6495	0.5162
w/ DARE (drop 90%)	0	0.4908	0.3162	0.0422	0.3682	0.3185	0.5114	0.4729
Ties-Merging [84]	0.2048	0.8440	0.8113	0.5819	0.8570	0.6465	0.7481	0.4296
w/ DARE (drop 30%)	0	0.5103	0.3382	-0.0024	0.3961	0.3238	0.5277	0.4838
w/ DARE (drop 50%)	0.0464	0.6021	0.5343	0.0192	0.6846	0.3410	0.5841	0.4982
w/ DARE (drop 70%)	0.1342	0.7833	0.7672	0.1667	0.8180	0.4172	0.691	0.5271
w/ DARE (drop 90%)	0.2618	0.8383	0.8039	0.6082	0.8336	0.5551	0.7692	0.5235

- **AdaMerging** [85] uses an unsupervised method to learn the merging coefficients for each task vector (Task-wise AdaMerging) or each layer (Layer-wise AdaMerging). AdaMerging++ is realized by adopting Ties-Merging [84] before learning the merging coefficients.
- **DARE** [90] (Drop and Rescale) validates the extremely redundant properties of language models. As a pre-processing technique, DARE randomly drops most (90% or even 99%) delta parameters (task vectors) before merging to potentially mitigate the interference of parameters among models.

## D More experimental results

### D.1 Merging ViT-B/16 models on 8 tasks

We follow the settings in Section 4.1.1 and merge ViT-B/16 models. Tab. 11 shows the accuracy of merging ViT-B/16 models on eight vision tasks. The proposed EMR-MERGING brings about 5.6% performance improvement compared to Adamerging++ [85], further demonstrating the effectiveness of EMR-MERGING.

### D.2 Merging ViT-B/32 models on 9 tasks (ImageNet-1K added)

To further explore the performance of EMR-MERGING, we follow the settings in Section 4.1.1 and add one more task, ImageNet-1K [18]. We merge models on these nine tasks using different merging methods. The results are shown in Tab. 12 and EMR-Merging shows a much more significant improvement compared to existing merging methods (up to 20%).

### D.3 DARE’s experimental results and causes

DARE’s experimental results when combined with RegMean and Task Arithmetic are shown in Tab. 13. It can be seen that when applied to merge eight models, DARE works on a few tasks under low dropping rate settings but it generally fails. We attribute its failure to the random dropping strategy’s unapplicability to merging multiple models. Under the setting of merging two or three models, randomly dropping most parameters in task vectors can significantly reduce interference but conflicts are a lot more difficult to avoid when merging multiple models.

Table 14: Performance of Task Arithmetic [30], Ties-Merging [84], Ties-Merging [84] w/ DARE [90], and RegMean [33] under different hyper-parameter settings.  $\lambda$  for task vector-based methods is the merging coefficient.  $P$  is the drop rate for DARE.  $a$  is the non-diagonal multiplier for RegMean.

Methods	Single-Sentence Tasks		Similarity and Paraphrase Tasks			Inference Tasks		
	CoLA	SST2	MRPC	STSB	QQP	MNLI	QNLI	RTE
Individual	0.6018	0.9404	0.8922	0.9063	0.9141	0.872	0.9271	0.7906
<b>EMR-MERGING (Ours)</b>								
	0.3996	0.9335	0.8627	0.8277	0.8972	0.8545	0.8957	0.7437
<b>Task Arithmetic</b>								
$\lambda = 0.1$	0.0464	0.742	0.6691	0.2344	0.771	0.3567	0.6919	0.556
$\lambda = 0.3$	0.1878	0.8589	0.799	0.7403	0.8378	0.5908	0.6967	0.6209
$\lambda = 0.5$	-0.0089	0.7913	0.7794	0.5686	0.8271	0.4631	0.5387	0.4693
$\lambda = 0.7$	-0.0079	0.6525	0.7819	0.1292	0.8146	0.3949	0.5279	0.5054
$\lambda = 0.9$	-0.0207	0.7202	0.4167	-0.1283	0.8012	0.2913	0.5294	0.5162
$\lambda = 1.0$	0	0.5619	0.3554	-0.2496	0.7939	0.259	0.5338	0.5162
<b>Ties-Merging</b>								
$\lambda = 0.1$	0	0.4908	0.3162	0.0214	0.3682	0.3186	0.5105	0.4729
$\lambda = 0.3$	0	0.5631	0.5049	-0.0074	0.4696	0.35	0.5649	0.4621
$\lambda = 0.5$	0.2232	0.7592	0.7696	0.1149	0.827	0.4486	0.6939	0.4368
$\lambda = 0.7$	0.2507	0.8291	0.7917	0.3774	0.8488	0.5858	0.7507	0.4188
$\lambda = 0.9$	0.2048	0.844	0.8113	0.5819	0.857	0.6465	0.7481	0.4296
$\lambda = 1.0$	0.1712	0.8406	0.799	0.6444	0.859	0.6409	0.7069	0.426
<b>Ties-Merging w/ DARE</b>								
$\lambda = 0.2, P = 0.3$	0	0.4920	0.3162	0.0053	0.3682	0.3186	0.5131	0.4477
$\lambda = 0.2, P = 0.5$	0	0.0043	0.3162	0.0036	0.3690	0.3202	0.5226	0.4946
$\lambda = 0.2, P = 0.7$	0.0464	0.6388	0.5735	0.0301	0.0047	0.3383	0.5984	0.5090
$\lambda = 0.2, P = 0.9$	0.2402	0.8165	0.7843	0.2696	0.8112	0.4384	0.7223	0.5415
$\lambda = 0.3, P = 0.3$	0	0.5103	0.3382	-0.0024	0.3961	0.3238	0.5277	0.4838
$\lambda = 0.3, P = 0.5$	0.0464	0.6021	0.5343	0.0192	0.6846	0.3410	0.5841	0.4982
$\lambda = 0.3, P = 0.7$	0.1342	0.7833	0.7672	0.1667	0.8180	0.4172	0.691	0.5271
$\lambda = 0.3, P = 0.9$	0.2618	0.8383	0.8039	0.6082	0.8336	0.5551	0.7692	0.5235
$\lambda = 0.4, P = 0.3$	0.0656	0.6216	0.5588	0.0192	0.7301	0.3461	0.5891	0.5162
$\lambda = 0.4, P = 0.5$	0.1172	0.7374	0.7451	0.1045	0.8157	0.3913	0.6667	0.5126
$\lambda = 0.4, P = 0.7$	0.2440	0.8234	0.7843	0.3955	0.8371	0.5496	0.7216	0.4838
$\lambda = 0.4, P = 0.9$	0.1380	0.8440	0.8064	0.7044	0.8365	0.5835	0.6529	0.5054
<b>RegMean</b>								
$a = 0.7$	0.3005	0.9037	0.7525	0.6349	0.8322	0.6794	0.8157	0.5632
$a = 0.8$	0.3346	0.9014	0.7549	0.6375	0.8339	0.6841	0.8173	0.5704
$a = 0.9$	0.3445	0.9048	0.7525	0.6362	0.8361	0.6918	0.821	0.5632
$a = 1.0$	0.3667	0.906	0.7574	0.6268	0.8355	0.7002	0.8235	0.5848

#### D.4 Results under different hyper-parameter settings

In Section 4.2.1, we presented the best performance of Ties-Merging, Task Arithmetic, and RegMean among multiple hyper-parameter settings. Here we present more experimental results of Ties-Merging, Task Arithmetic, and RegMean under different hyper-parameter settings in Tab. 14.

#### D.5 Detailed information for merging different number of models

In Section 4.4, we showed partial results of merging different number of ViT-B/32 models by Fig. 6. Here we provide quantified and task-specific performance results in Tab. 15.

Table 15: Merging different number of ViT-B/32 models.

Methods	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD	Avg Acc
<b>Individual</b>									
2 Tasks	75.3	77.7	-	-	-	-	-	-	76.5
3 Tasks	75.3	77.7	96.1	-	-	-	-	-	83.0
4 Tasks	75.3	77.7	96.1	99.7	-	-	-	-	87.2
5 Tasks	75.3	77.7	96.1	99.7	97.5	-	-	-	89.3
6 Tasks	75.3	77.7	96.1	99.7	97.5	98.7	-	-	90.8
7 Tasks	75.3	77.7	96.1	99.7	97.5	98.7	99.7	-	92.1
8 Tasks	75.3	77.7	96.1	99.7	97.5	98.7	99.7	79.4	90.5
<b>Ties-Merging</b>									
2 Tasks	69.2	68.2	-	-	-	-	-	-	68.7
3 Tasks	69.2	68.0	78.9	-	-	-	-	-	72.0
4 Tasks	68.9	67.9	79.4	86.0	-	-	-	-	75.5
5 Tasks	68.6	67.1	79.0	83.5	66.6	-	-	-	73.0
6 Tasks	68.0	66.4	77.9	80.1	74.4	69.9	-	-	72.8
7 Tasks	66.6	65.7	75.7	76.7	81.0	69.2	96.4	-	75.9
8 Tasks	64.8	62.9	74.3	78.9	83.1	71.4	97.6	56.2	72.4
<b>EMR-MERGING (Ours)</b>									
2 Tasks	78.9	76.1	-	-	-	-	-	-	77.5
3 Tasks	77.9	75.2	95.3	-	-	-	-	-	82.8
4 Tasks	77.4	74.9	94.8	99.7	-	-	-	-	86.7
5 Tasks	77.2	74.2	94.7	99.7	97.1	-	-	-	88.6
6 Tasks	76.4	73.4	94.2	99.7	97.0	98.5	-	-	89.9
7 Tasks	75.8	73.3	93.6	99.6	96.9	98.2	99.6	-	91.0
8 Tasks	75.2	72.8	93.5	99.5	96.9	98.1	99.6	74.4	88.7

Table 16: Sparsity (ratio of non-zero items) of the masks and the values of the rescalers when merging ViTs on 8 vision tasks and RoBERTa models on 8 language tasks.

<b>Sparsity</b>	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD
ViT-B/32	0.7194	0.7121	0.7106	0.6994	0.7195	0.7062	0.7132	0.7058
ViT-L/14	0.6832	0.6699	0.6734	0.6579	0.6748	0.6444	0.6614	0.6620
<b>Rescalers</b>	SUN397	Cars	RESISC45	EuroSAT	SVHN	GTSRB	MNIST	DTD
ViT-B/32	0.7489	0.7635	0.7489	0.7476	0.7962	0.7652	0.7981	0.7624
ViT-L/14	0.7656	0.7652	0.7537	0.7384	0.7874	0.7313	0.7763	0.7638
<b>Sparsity</b>	CoLA	SST2	MRPC	STSB	QQP	MNLI	QNLI	RTE
RoBERTa	0.6264	0.6547	0.6498	0.6150	0.7620	0.7739	0.6243	0.5979
<b>Rescalers</b>	CoLA	SST2	MRPC	STSB	QQP	MNLI	QNLI	RTE
RoBERTa	0.2458	0.4698	0.5033	0.2078	0.8891	0.8987	0.4683	0.1466

## D.6 Sparsity of masks and values of rescalers.

We show the sparsity of the masks and the values of the rescalers when merging eight ViTs and eight RoBERTa models in Tab. 16.

## E More visualization results

In Section 3, we showed some visualization results using t-SNE [69] and Grad-CAM [61]. Here we provide more visualization results of both existing merging methods and EMR-MERGING. t-SNE and Grad-CAM visualization results are shown in Fig. 8 and Fig. 9, respectively.

## **F Configuration of Fig. 4 and Fig. 7**

In Fig. 4 and Fig. 7, we hope to compare the sign conflicts, L2 distance, and cosine similarity of the merged model weights and individual model weights. To calculate the sign conflicts, we element-wisely compare the merged model weights to each individual model weights and record the ratio of the elements whose signs conflict. We report the average value of the sign conflicts between the merged model and each individual model. To calculate the L2 distance or cosine similarity, we first flatten the merged model weights and each individual model weights as 1-dimension vectors. Then we calculate the L2 distance or cosine similarity between the merged model and each individual model and report the average value.

## **G Limitations and future works**

Despite the convincing results, the proposed method suffers from several limitations. On the one hand, compared to existing methods, EMR-MERGING requires a little additional memory to store the light-weight task-specific modulators. On the other hand, as a common limitation of task vector-based methods, EMR-MERGING cannot be generalized to models trained from-scratch because the task vector is based on the pretrain-finetune paradigm.

Further improving the performance of the merged model and generalizing model merging to models trained from-scratch or even models with different structures are significant directions for future work. Additionally, combining model merging with low bit-width quantization has broad application prospects and is also a potential future work.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in both abstract and Section 1 accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Please check Appendix G.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]



Justification: The full set of assumptions and a complete (and correct) proof are detailed in Appendix B.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We disclose the information in Section 4 and we provide the code and data for the convenience of reproduction.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the code for reproduction. Please check Abstract.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We specify them in Section 4 and in our released code.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Error bars are not reported.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: The required computer resources are decided by the structure of the models to be merged.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The research is conducted with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Not applicable to societal impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original papers or websites that produced the code package or dataset.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.