

Model Stock: All we need is just a few fine-tuned models

Dong-Hwan Jang^{1,2*} Sangdoo Yun^{1†} Dongyoon Han^{1†}

¹ NAVER AI Lab ² Samsung Advanced Institute of Technology (SAIT)

[†] corresponding authors

Abstract. This paper introduces an efficient fine-tuning method for large pre-trained models, offering strong in-distribution (ID) and out-of-distribution (OOD) performance. Breaking away from traditional practices that need a multitude of fine-tuned models for averaging, our approach employs significantly fewer models to achieve final weights yet yield superior accuracy. Drawing from key insights in the weight space of fine-tuned weights, we uncover a strong link between the performance and proximity to the center of weight space. Based on this, we introduce a method that approximates a center-close weight using only two fine-tuned models, applicable during or after training. Our innovative layer-wise weight averaging technique surpasses state-of-the-art model methods such as Model Soup, utilizing only two fine-tuned models. This strategy can be aptly coined *Model Stock*, highlighting its reliance on selecting a minimal number of models to draw a more optimized-averaged model. We demonstrate the efficacy of Model Stock with fine-tuned models based upon pre-trained CLIP architectures, achieving remarkable performance on both ID and OOD tasks on the standard benchmarks, all while barely bringing extra computational demands. Our code and pre-trained models are available at <https://github.com/naver-ai/model-stock>.

1 Introduction

Pre-train/fine-tune paradigm [16, 19, 33, 35, 36] has proven to be a strong framework for training models to reach state-of-the-art performance. This approach, especially pivotal in fine-tuning pre-trained models, involves models acquiring general knowledge during pre-training and task-specific knowledge during fine-tuning. How we perform a fine-tuning stage is crucial, affecting task performance and robustness against distribution shifts.

Recent advancements, notably Model Soup [35], which merges weights from multiple fine-tuned models trained under different training setups, have shown impressive performance without increasing inference costs. This method is believed to be effective because these models often reside in the same loss basin, and their merging results in a lower and flat loss basin. However, Model Soup’s requirement for multiple fine-tuned models (more than dozens) raises concerns

* Work done during an internship at NAVER AI Lab.




Fig. 1: Model Stock vs. Model Soup. Model Stock consistently enjoys improved accuracy on ImageNet (x-axis) and distribution shift benchmarks (y-axis) against individual fine-tuned models (gray circles). We plot WiSE-FT [36] curves to each fine-tuned model, highlighting Model Stock’s better performance on distribution shifts compared to Model Soup [35]. Note that Model Stock has much smaller computational costs than Model Soups ($24\times$ smaller), *i.e.*, Model Stock requires two fine-tuning procedures, whereas Model Soups are leveraging 48 various fine-tuned models in this experiment.

about efficiency and practicality in general scenarios where the models need to be prepared from scratch. Thus, our question is: Is there an efficient way to achieve an effective merged weight from very few fine-tuned models?

We initially explore the dynamics of fine-tuned weights under the simplest scenario: varying the random seeds while maintaining the other training setups. It reveals that the fine-tuned weights with different random seeds reside on a very thin shell layer-wise during and after training. We then delve into the impact of a model-soup-like weight averaging approach. Our findings show that the closer proximity of the averaged weights correlates with improved In-Distribution (ID) and Out-Of-Distribution (OOD) performance.

Building upon the findings, we propose a novel approach of fine-tuning method coined **Model Stock**, analogous to chicken stock in cooking, distinguishing it from what Model Soup intended. Now, the answer to our question is indeed affirmative: Model Stock approximates the merged weight using just a few fine-tuned models, leveraging the weight space’s geometric properties and a pre-trained model’s anchoring effect. This strategy offers a more computationally efficient alternative to the labor-intensive averaging of fine-tuned models, streamlining the process while enhancing model performance. Fig. 1 illustrates our brief comparison of Model Stock vs. Model Soup [35]. We reproduce the Model Soup experiments¹ based on the CLIP ViT-B/32 initialization by fine-tuning 48 models with various hyper-parameters, which is defined as *zero-shot* initialization

¹ We follow the standard grid hyper-parameter sweep [35], fine-tuning for 10 epochs. Our results align with those presented in Fig. D.1 (right) of the original paper [35].

setting. Fig. 1 shows that Model Stock outperforms Model Soups with much smaller computational costs.

Our comprehensive experiments demonstrate the effectiveness of Model Stock. We achieve performance comparable to, or even surpassing, that of the more resource-intensive methods such as Model Soup [35], using only a fraction of the models. Specifically, our method achieves 87.8% ImageNet top-1 accuracy (ID) and averaged 74.9% in five distribution shift benchmarks (OOD) on ViT-L/14 fairly compared with the prior arts using the CLIP pre-trained weight. We believe that our study not only underscores Model Stock’s practicability but also opens new directions in the pre-train/fine-tune paradigm for superior performance across various tasks.

2 Analyzing Fine-tuned Weights

Our study is driven by two fundamental findings related to the performance and robustness of fine-tuned models. The first one is that model weights are fine-tuned on different random seeds² lie on *thin shell* in weight space layer-wise. The second posits that closer proximity to the center of this thin shell is beneficial for improving performance across the ImageNet and distribution shift benchmarks. The substantiation and implications of these observations are discussed in the subsequent sections.

2.1 Geometric Properties Between Weights

Angle and norm of weights. We begin by examining the intrinsic properties of the weights in fine-tuned models. We define the weight vector of the fine-tuned model at k -th layer as $\mathbf{w}^{(k)} \in \mathbb{R}^{n^{(k)}}$ where $n^{(k)}$ is the number of weight parameters at k -th layer, and the origin $\mathbf{0}$ as the pre-trained model weight $\mathbf{w}_0^{(k)}$ at k -th layer. Then the angle $\theta^{(k)}$ between two weight $\mathbf{w}_1^{(k)}$ and $\mathbf{w}_2^{(k)}$ is defined by: $\theta^{(k)} = \arccos \left(\frac{\mathbf{w}_1^{(k)} \cdot \mathbf{w}_2^{(k)}}{\|\mathbf{w}_1^{(k)}\| \|\mathbf{w}_2^{(k)}\|} \right)$, where the Euclidean l_2 -norm $\|\mathbf{w}^{(k)}\|$ of a $n^{(k)}$ defined as $\|\mathbf{w}^{(k)}\| = \sqrt{\sum_{i=1}^{n^{(k)}} w_i^{(k)2}}$. Angle and norm will provide a geometric view of the weights at k -th layer between fine-tuned models.

Observation 1: Angle and norm consistency among fine-tuned weights.

We investigate the weight space of models fine-tuned on ImageNet from a pre-trained model with various random seeds. Our first observation is that both angle $\theta^{(k)}$ between two different models and norm $\|\mathbf{w}^{(k)}\|$ of a weight exhibit *consistent* values with very low standard deviations, as shown in Fig. 2. This consistency can be mathematically represented as follows: For all i and $j \in [1, N]$ when the number of fine-tuned weights, N , is sufficiently large, the following holds:

$$\mathbf{w}_i^{(k)} \cdot \mathbf{w}_j^{(k)} = \begin{cases} (l^{(k)})^2 & \text{if } i = j, \\ (l^{(k)})^2 \cos \theta^{(k)} & \text{otherwise,} \end{cases} \quad (1)$$

² Random seed influences training randomness, such as training data shuffling and data augmentation parameters.

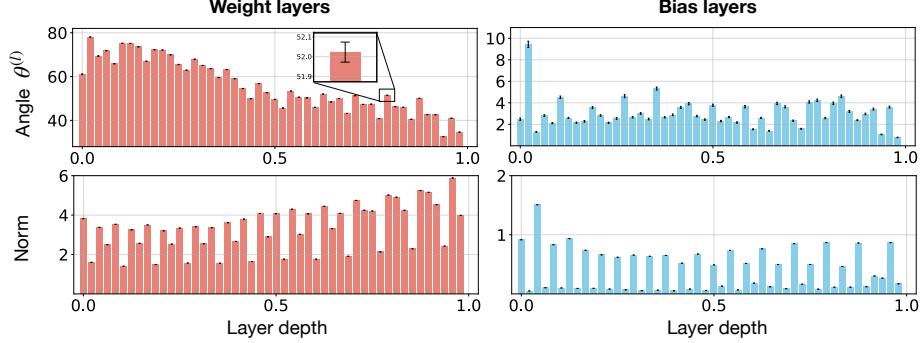


Fig. 2: Layer-wise angle and norm of fine-tuned models. We measure the angle $\theta^{(k)}$ (degree) and norm $\frac{\|\mathbf{w}^{(k)}\|}{\sqrt{n^{(k)}}}$ for 50 distinct weights, fine-tuned under different random seeds. We separately visualize weight layers (red bars) and bias layers (blue bars), where bias layers have much smaller angles. We plot the mean angle and norm values with standard deviation (black error bar). The results show that any two fine-tuned weights have *layer-wise consistent* angle and norm with extremely low standard deviation.

where $l^{(k)}$ and $\theta^{(k)}$ are constants that describe the magnitude and angle between weights at k -th layer, respectively. Henceforth, to simplify notation, we will omit the superscript (k) indicating the layer index.

Interestingly, these consistencies in angle and norm are observed 1) across diverse setups and 2) both during and after training. Fig. 2 shows this consistency over 50 fine-tuned CLIP ViT-B/32³ [26] models. It illustrates that the *layer-wise* norm and angle of these models exhibit almost constant values with extremely minimal error. While this figure depicts a specific model (*i.e.*, CLIP ViT-B/32), we establish that such regularity is not confined to a single model or setting but is consistent across various CLIP fine-tuning scenarios. We conjecture this holds irrespective of networks (ViT [5], Hybrid-ViT [5], ResNet [8], ConvNext [20]), optimizers (SGD, AdamW [21]), data augmentations (RRC [31], RandAug [4]), datasets (CIFAR [15], ImageNet [30]), or initialization of the classifier (zero-shot, LP-FT [16]). Remarkably, this regularity also holds for fine-tuned weights at each step during training as well as after training. A comprehensive analysis supporting these findings is presented in the Appendix.

Based on the observation, we presume the distribution of the fine-tuned weights. The center of the fine-tuned weights is defined as $\boldsymbol{\mu} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{w}_i$. We then deduce the following properties among fine-tuned weights: (i) $\|\mathbf{w}_i - \boldsymbol{\mu}\| = \text{constant}$, indicating a *thin shell* distribution; (ii) $(\mathbf{w}_0 - \boldsymbol{\mu}) \perp (\mathbf{w}_i - \boldsymbol{\mu})$; and (iii) $(\mathbf{w}_i - \boldsymbol{\mu}) \perp (\mathbf{w}_j - \boldsymbol{\mu})$ for all $i, j \in [1, N]$. These properties are depicted in Fig. 5 for better understanding. The detailed proof is in the Appendix.

³ <https://github.com/openai/CLIP>

Table 1: Distance from the center (i.e., $\|\mathbf{w} - \mu\|$) vs. performance. We report the ImageNet and distribution shift performance with the distance from the center of weights μ for fine-tuned and averaged models. We observe 1) both models consistently maintain a nearly constant distance from μ with remarkably small standard deviation; 2) averaging more models approaches μ , boosting ID and OOD performance.

	$\ \mathbf{w} - \mu\ $	ImageNet	Avg. shifts
Fine-tuned	$13.133 \pm .004$	79.72	46.37
$\mathbf{w}_{\text{avr}}^{(2)}$	$9.192 \pm .003$	80.24	47.76
$\mathbf{w}_{\text{avr}}^{(3)}$	$7.439 \pm .025$	80.35	48.18
$\mathbf{w}_{\text{avr}}^{(5)}$	$5.633 \pm .014$	80.47	48.53
$\mathbf{w}_{\text{avr}}^{(50)} \simeq \mu$	~ 0	80.59	48.85




Fig. 3: Test error landscape. We visualize a test error landscape with three weight points: a pre-trained model (\mathbf{w}_0), a fine-tuned model (\mathbf{w}_1), and the averaged weights of 50 fine-tuned models ($\mathbf{w}_{\text{avr}}^{50}$). $\mathbf{w}_{\text{avr}}^{50}$ locates near the lowest error basin. A better solution (\mathbf{w}_H) with lower error than \mathbf{w}_1 can be easily found utilizing those three weights (§3).

2.2 Center of Weights and Performance

We proceed to explore the relationship between the proximity to the center of fine-tuned weights and their performance on ID and OOD datasets. Given that computing the exact center is infeasible, we approximate it by averaging differently seeded 50 fine-tuned weights, using it as a pseudo-center (i.e. $\mu \simeq \mathbf{w}_{\text{avr}}^{(50)}$).

Observation 2: Distance from the center of weights and performance.

Table 1 offers quantitative observations about the fine-tuned weights and their performance using CLIP ViT-B/32. The results include distances from the weight center (μ) of fine-tuned models and their averaged counterparts with their ID and OOD performances. Going closer to the center by averaging the weights leads to improving both performances. Interestingly, the standard deviation of the distances is less than 0.1% of the mean distance, suggesting highly consistent fine-tuned weight distances from the center across different weights. This suggests that fine-tuned weights occupy a *thin shell* as discussed in §2.1.

Observation 3: Fine-tuned weights occupy local minima edges in the test error landscape.

We present an additional observation regarding the test error landscape, which relates to the performance around the center of weights. Fig. 3 depicts the test error landscape on the ImageNet test dataset within a two-dimensional plane. This plane includes a pre-trained model’s weight (\mathbf{w}_0), a single fine-tuned model (\mathbf{w}_1), and the pseudo-center ($\mathbf{w}_{\text{avr}}^{(50)}$). This landscape




Fig. 4: Weights closer to the center enjoy higher performances. We show the ImageNet accuracy of weights with their distance from the center (μ), which is approximated with $\mathbf{w}_{\text{avr}}^{(50)}$. Averaged weights are closer to the center than individual weights, and their accuracy increases as the averaging number of models (N) increases. Gray circles are the weights randomly sampled from the Gaussian distribution centered at $\mathbf{w}_{\text{avr}}^{(50)}$. Even the random weights also achieve higher performance as they reach the center. The results indicate the critical role of proximity to the center on performance.

reveals that a fine-tuned model typically occupies the boundary of test error regions. On the other hand, centered near pseudo-center ($\mathbf{w}_{\text{avr}}^{(50)}$), the test error is the lowest and gets higher as the weights get far from the center. Interestingly, along the line segment $\overline{\mathbf{w}_0\mathbf{w}_1}$ (i.e., a WiSE-FT curve [36]), the fine-tuned weight \mathbf{w}_1 is neither the point closest to the pseudo-center nor the one with the lowest test error. We will connect this observation in §3 to find \mathbf{w}_H , the weight on the line closest to the center.

Observation 4: Randomly perturbed weights nearing the center also merit high performance. To further investigate the impact of proximity to the center on performance, we conduct a toy experiment to measure ImageNet performance using random weights generated by adding layer-wise noise to the weight at the center. The standard deviation of the noise for each layer is adjusted to align with the distribution of fine-tuned weights. Fig. 4 presents a scatter plot correlating the distance of model weights from the distribution’s center with corresponding ImageNet accuracies. Remarkably, randomly sampled weights demonstrate performance comparable to fine-tuned and averaged models, highlighting the importance of center proximity.

Finally, the above observations naturally raise a question: Why do fine-tuned weights through optimization not reach the center, staying constantly close instead? Previous studies [3, 13] tell us that optimization steps might struggle to guide fine-tuned weights to the center of the weight distribution due to the many stationary points in the loss surface. Alternatively, averaging independently fine-tuned models is a unique solution but both laborious and resource-intensive. It appears that there are no better alternatives for getting closer to the center, as




Fig. 5: Comprehensive illustration of the geometric dynamics of fine-tuned weights. This figure illustrates the behavior of fine-tuned weights that lie on a thin shell, supporting our Gaussian distribution hypothesis. Each sphere represents a thin shell that fine-tuned models lie on at each training step, with \mathbf{w}_0 denoting the pre-trained model. The curved lines trace the fine-tuning trajectory from \mathbf{w}_0 , while the red vectors indicate that fine-tuned models at each time step are equidistant from \mathbf{w}_0 .

optimization proves ineffective near these flat local minima. Could there be a faster approach to reaching the center? This question will be addressed in §3.

2.3 Our Hypothesis

The observed geometric patterns in weight distributions closely align with mathematical properties of Gaussian distributions, represented as $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$. Therefore, a plausible reason for the particular geometric pattern of fine-tuned weights could be the influence of Gaussian noise within the weight space. In high-dimensional spaces, vectors sampled from such distributions tend to have nearly identical norms, specifically $\sqrt{\|\boldsymbol{\mu}\|^2 + \text{trace}(\Sigma)}$ and consistent in-between angles, due to the *concentration of measure phenomenon* [17]. The likelihood of the squared norm significantly deviating from this expected value is exponentially negligible in high-dimensional spaces, like a weight space.

In other words, the vectors sampled from high-dimensional Gaussian distribution lie on a very thin shell with the radius $\approx \sqrt{\text{trace}(\Sigma)}$ around the center $\boldsymbol{\mu}$. Consequently, we hypothesize that fine-tuned weights follow a Gaussian distribution in a layer-wise manner. While not necessary for our observation, this hypothesis provides a sufficient condition for understanding the geometric dynamics of fine-tuned weights. For example, it aids in the intuitive understanding of the fact that the distance of $\mathbf{w}_{\text{avr}}^{(N)}$ from the weight center is proportional to $1/\sqrt{N}$, signifying the reduction of variance. Fig. 5 comprehensively illustrates the observations and our hypothesis discussed in §2.

3 Method

This section introduces our method, Model Stock, a cost-efficient weight merging method. As discussed in §2.2, getting closer to the center of weights $\boldsymbol{\mu}$ induces improved model performance. A straightforward method to approximate




Fig. 6: Schematic concept of our Model Stock. We present two scenarios with a small angle (left) and a large angle (right). Given a pre-trained weight (\mathbf{w}_0) and two fine-tuned weights (\mathbf{w}_1 , and \mathbf{w}_2), we visualize a gray triangle representing the span of three weights. We consider this triangle area our search space spanned by three weights. We aim to find the best weight point on the triangle nearest the ideal center μ . We find that the perpendicular foot \mathbf{w}_H from point μ to the plane is the nearest point, which can be specified solely using the angle between the fine-tuned models, even without knowing the exact position of the center μ . We utilize \mathbf{w}_H as the merged weight of Model Stock. Intuitively, when the angle θ is large (*i.e.*, two weights are diverse), as in the right figure, \mathbf{w}_H will rely more on the pre-trained weight (\mathbf{w}_0), vice versa.

μ is averaging multiple model weights, which can be computationally expensive. We propose an efficient alternative method by leveraging *the pre-trained model weights*, an aspect previously neglected by existing weight-merging methods. A pre-trained model usually possesses general knowledge and shows robust and reliable performance in out-of-distribution (OOD) cases [26]. Therefore, a pre-trained model can become a *robust anchor point*. As shown in Fig. 3, we could readily identify a weight (\mathbf{w}_H) that is closer to the center—and thus better—by interpolating with the anchor. Building on this concept, we propose a method to approximate the center of weights more accurately with only a few fine-tuned weights. Again, for readability, we omit the layer notation (k), but the subsequent method is applied layer-wise.

On two fine-tuned models. We observed in §2 that two fine-tuned models with different random seeds have almost constant norms and the angle between them. Based on this, we define a plane connecting the pre-trained model and two fine-tuned models as shown in Fig. 6 (depicted as a gray triangular area). Any weight vector on this plane can be expressed as a linear combination of the pre-trained model and two fine-tuned models. Our goal is to find the weight closest to the center of fine-tuned weights on this plane, which is the perpendicular foot (\mathbf{w}_H) from the center of distribution (μ) to the plane.

Even without knowing the exact position of the center μ , \mathbf{w}_H can be specified solely using the angle between the fine-tuned models with the following two conditions that μ must satisfy. First, as mentioned in §2.1, $(\mathbf{w}_1 - \mu) \perp (\mathbf{w}_2 - \mu)$ and $\|\mathbf{w}_1 - \mu\| = \|\mathbf{w}_2 - \mu\|$ hold, implying that $\triangle \mathbf{w}_1 \mu \mathbf{w}_2$ forms an isosceles right triangle. In Fig. 6, μ should lie on the dotted hyper-circle. Another condition is that the condition $(\mathbf{w}_0 - \mu) \perp (\mathbf{w}_{12} - \mu)$ must be satisfied, where $\mathbf{w}_{12} = \frac{\mathbf{w}_1 + \mathbf{w}_2}{2}$. This condition arises from the second property in §2.1, where $(\mathbf{w}_0 - \mu) \perp (\mathbf{w}_i - \mu)$. Combining these two conditions, μ is at the point where the line starting from \mathbf{w}_0

is tangent to the hyper-circle. We precisely determine the position of μ within a 3D volume slice that encompasses both μ and $\Delta \mathbf{w}_0 \mathbf{w}_1 \mathbf{w}_2$. Consequently, we can find the position of the closest weight \mathbf{w}_H to the distribution’s center on the plane using straightforward geometric principles. The position of the perpendicular foot is determined as follows:

$$\mathbf{w}_H = \frac{2 \cos \theta}{1 + \cos \theta} \cdot \mathbf{w}_{12} + \left(1 - \frac{2 \cos \theta}{1 + \cos \theta}\right) \cdot \mathbf{w}_0. \quad (2)$$

For more detailed proof, please refer to the Appendix. Note that the interpolation ratio $t = \frac{2 \cos \theta}{1 + \cos \theta}$ is solely determined by the angle θ between two fine-tuned models. Crucially, unlike previous methods, determining t does not require extra training [19, 32, 33, 35] or heuristic hyper-parameter settings [6, 35], thereby simplifying the process and enhancing its accessibility and efficiency.

As the θ decreases, the pre-trained model is less utilized for merging, as shown in Fig 6 (left). Coupled with the observation in Fig. 2, this indicates that bias layers rely less on pre-trained models and focus more on fine-tuned models, whereas weight layers depend more on pre-trained models. This observation extends the findings of previous works such as BitFit [38] and LP-FT [16]. In the case of BitFit and LP (*i.e.*, the first step of LP-FT), bias and classifier layers fully utilize fine-tuning, while other weight layers (attention and MLP) rely on pre-trained models. We present an additional analysis in the Appendix.

On N fine-tuned models. We further extend the previous derivation to N fine-tuned models to move even closer to the weight center. Let us denote $\mathbf{w}_{\text{avr}}^{(N)}$ as the N -averaged weight, $\sum_{i=1}^N \mathbf{w}_i / N$, and $\mathbf{w}_H^{(N)}$ as the weight in the span($\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_N$) closest to the distribution’s center. Then, we derive the position of $\mathbf{w}_H^{(N)}$ as:

$$\mathbf{w}_H^{(N)} = t \cdot \mathbf{w}_{\text{avr}}^{(N)} + (1 - t) \cdot \mathbf{w}_0, \quad \text{s.t.} \quad t = \frac{N \cos \theta}{1 + (N - 1) \cos \theta}. \quad (3)$$

Detailed proof is in the Appendix. Similar to the case of two fine-tuned models, the interpolation ratio t depends solely on the angle θ between the pre-trained model and the N fine-tuned models.

Periodic merging. To move one step forward with our method, we propose *periodic merging*, which is performed between the fine-tuned models and the pre-trained model during training. As the geometric properties of weights are also applicable to weights during training (refer to Appendix for more details), we fully utilize this phenomenon here.

We strategically merge weights at the end of every epoch, allowing for the fine-tuning process to be parallelized with merging. The interpolation ratio for merging is determined by the angle θ between the pre-trained model and the fine-tuned models at the current epoch. Fig. 7 visualizes this periodic merging process. We argue that employing the periodic merging can approximate the




Fig. 7: Periodic merging for Model Stock. Weights are merged every epoch.

center of weights more accurately. In §4.3, we empirically show that the periodic merging yields superior performance and achieves a closer distance to the center.

4 Experiment

We present the key experimental results in this section. We first provide our experimental setups in §4.1. Then, we present the main results in §4.2, and ablation studies in §4.3. More detailed experimental setups, analysis, and further results are in the Appendix. We will release our codes and weights publicly.

4.1 Experimental Setup

Models. We conduct the experiments on CLIP ViT-B/32, CLIP ViT-B/16, and CLIP ViT-L/14 models. We set the number of fine-tuning models for Model Stock as two. We compare Model Stock against various fine-tuning techniques, including Model Soups [35], LP-FT [16], CAR-FT [23], FTP [33], and FLYP [7]. We use CLIP ViT-B/32 for ablation studies.

Datasets. We fine-tune models on the ImageNet-1K [30] training dataset. We report the ImageNet-1K [30] top-1 accuracy for evaluating in-distribution (ID) performance. For distribution shift scenarios, we consider five out-of-distribution (OOD) benchmarks including ImageNet-V2 [29], ImageNet-R [9], ImageNet-Sketch [34], ImageNet-A [10], and ObjectNet [1]. Since previous methods, except Model Soups [35], have not been evaluated on ObjectNet, we omit the ObjectNet results when comparing against them.

Training setup. We initialize the classifier weights (*e.g.*, 1000 classes for ImageNet) using the text encoder of CLIP and text prompts following Model Soup [35]. We use AdamW [21] with batch size 512 and weight decay 0.1 for all the experiments, including vanilla fine-tuning, Model Soup reproduction, and Model Stock⁴. We employ two training setups for comparisons: (1) training 10 epochs with minimal data augmentation, following Model Soup’s zero-shot initialization setup, and (2) training 16 epochs with strong data augmentation, following one of Model Soup’s LP initialization setup [35], denoting with *. These setups enable a balanced comparison of Model Stock against zero-shot and LP-initialized Model Soups. Detailed hyper-parameters are in the Appendix.

4.2 Main Results

CLIP ViT-B/32. Table 2 shows the results of Model Stock on the pre-trained CLIP ViT-B/32 model by comparing it with Model Soups. ‘Avg. shifts’ denotes the average accuracy of the five OOD benchmark scores. Our Model Stock and Model Stock* show competitive performance with Model Soups. Furthermore, Model Stock* achieves state-of-the-art performance on ImageNet with **81.19%** top-1 accuracy. As described in Fig. 1, Model Stock with WiSE-FT [36] enjoys a

⁴ We reduce the batch size to 64 on CLIP ViT-L/14 due to memory limitation.

Table 2: Comparison against Model Soups [35] on CLIP ViT-B/32. We report the performance and relative fine-tuning costs on the CLIP ViT-B/32 scenario. α denotes the cost for LP initialization. Model Stock shows comparable performance with Model Soups with significantly reduced training costs.

Method	ImageNet	Avg. shifts	Cost
<i>Comparing with Model Soups from zero-shot init.</i>			
CLIP zero-shot Initialization	63.34	48.51	0
Vanilla FT	78.35	47.03	1
Uniform Model Soup (from zero-shot)	79.76	52.08	48
Greedy Model Soup (from zero-shot)	80.42	50.83	48
Model Stock	<u>79.89</u>	<u>50.99</u>	2
<i>Comparing with Model Soups from LP init.</i>			
CLIP LP initialization	75.57	47.21	α
Vanilla FT*	79.72	46.37	1
Uniform Model Soup (from LP init)	79.97	51.45	71+ α
Greedy Model Soup (from LP init)	<u>81.03</u>	<u>50.75</u>	71+ α
Model Stock*	81.19	48.69	2

Table 3: Model Stock on CLIP ViT-B/16. Model Stock shows competitive performance against previous fine-tuning methods on ImageNet and distribution shifts.

Method	ImageNet	Distribution shifts				
		Avg. shifts	IN-V2	IN-R	IN-A	IN-Sketch
Zero-shot	68.3	59.5	62.0	77.7	<u>49.9</u>	48.3
Vanilla FT	82.8	57.7	72.9	66.4	43.7	48.0
Vanilla FT*	83.7	57.4	73.5	67.6	40.0	48.6
LP [16]	79.7	48.1	71.5	52.4	27.8	40.5
LP-FT [16]	81.7	<u>60.5</u>	71.6	<u>72.9</u>	49.1	48.4
CAR-FT [23]	83.2	59.4	73.0	71.3	43.7	49.5
FTP [33]	<u>84.2</u>	49.7	74.6	47.2	26.5	50.2
FLYP [7]	82.6	<u>60.5</u>	73.0	71.4	48.1	49.6
Model Stock	84.1	62.4	<u>74.8</u>	71.8	51.2	51.8
Model Stock*	85.2	60.1	75.3	68.7	45.0	<u>51.3</u>

superior ID-OOD performance curve compared to Model Soup and its WiSE-FT curves. Note that Model Soups require dozens of fine-tuned models (*e.g.*, zero-shot and LP-init Model Soups use 48 and 71 models, respectively), highlighting the effectiveness of Model Stock along with efficiency utilizing only two models. We provide further comparison results with WiSE-FT [36] curves on LP-init Model Soups in the Appendix.

CLIP ViT-B/16. Table 3 presents a comprehensive comparison of different fine-tuning methods applied to CLIP ViT-B/16. Previous works [16, 32, 33] lack ObjectNet [1] results; therefore, we omit ObjectNet and report the other four OOD benchmarks. Complete results with ObjectNet and ImageNet-Real [2] are in the Appendix. The results show Model Stock exhibits exceptional performance on the ImageNet accuracy, *e.g.*, Model Stock* achieves 85.2% top-1 accuracy

Table 5: Impact of the number of fine-tuning models (N) on Model Stock. IN denote ImageNet accuracy.

	IN	Avg. Shifts	$\ \mathbf{w} - \boldsymbol{\mu}\ $
FT	79.7	46.7	13.13
$N=2$	80.1	48.8	10.01
$N=3$	80.2	48.8	9.05
$N=4$	80.4	48.9	8.45

Table 6: Impact of merging period on Model Stock. IN denotes ImageNet accuracy.

Period	IN	Avg. Shifts
1000 iters	79.8	48.7
5000 iters	79.9	48.5
1 epoch	80.1	48.8

on ImageNet, which is a state-of-the-art level. Model Stock also shows robust performance across diverse distribution shift scenarios.

CLIP ViT-L/14. Table 4 shows the results of Model Stock on the CLIP ViT-L/14 model. The results show that Model Stock can push the limit of benchmark scores with large-size backbone architecture. We remark that Model Stock* achieves state-of-the-art performance with 87.7% ImageNet top-1 accuracy, implying that Model Stock is still effective in a scale-up scenario. The results consistently demonstrate the high efficacy and robustness of Model Stock across diverse scales of models and various benchmark scenarios, reaffirming its potential in practical applications.

Table 4: Model Stock on CLIP ViT-L/14.

	IN	Avg. shifts
Zero-shot	75.0	63.0
Vanilla FT	85.8	66.8
Vanilla FT*	87.1	68.0
TPGM [32]	87.0	69.4
CAR-FT [23]	87.1	67.8
Model Stock	87.0	71.6
Model Stock*	87.7	73.5

4.3 Ablation studies and analysis of Model Stock

We conduct ablation studies on CLIP ViT-B/32. We train vanilla fine-tuned models and Model Stock for 16 and 8 epochs, respectively; thus, the training cost of Model Stock matches with a single fine-tuning process.

Experiments on the number of fine-tuned model N . Table 5 shows the effect of the number of fine-tuned models. The results show that Model Stock obtains enhanced performance and closer distance from the (pseudo-) center ($\|\mathbf{w} - \boldsymbol{\mu}\|$) as the number of merging models increases. Considering the trade-off between the performance and training cost induced by increased N , our setting ($N=2$) shows the best for Model Stock.

Study on the merging period of Model Stock. Table 6 shows the results of various merging periods, including 1000 and 5000 iterations settings. Note that 1 epoch is ~ 2500 iterations in our experiment. Model Stock shows consistent performance with various periods.

The post-training merging strategy of Model Stock. We study an alternative of Model Stock that merges fine-tuned weights only once after each fine-tuning process is finished, similar to Model Soups [35]. We denote it as Model

Table 7: Post-training merging strategy of Model Stock. We present ImageNet accuracy, distribution shifts, and distance from the center with the results of uniform averaging, a straightforward baseline.

	Uniform averaging ($\mathbf{w}_{\text{avg}}^N$)			Model Stock (post-training)		
	ImageNet	Avg. Shifts	$\ \mathbf{w} - \boldsymbol{\mu}\ $	ImageNet	Avg. Shifts	$\ \mathbf{w} - \boldsymbol{\mu}\ $
$N=2$	80.2	47.8	9.19	80.3(+0.1)	50.4(+2.6)	7.62(-1.57)
$N=3$	80.4	48.2	7.44	80.4(+0.0)	50.2(+2.0)	6.49(-0.95)
$N=4$	80.5	48.5	5.63	80.5(+0.0)	49.8(+1.4)	5.16(-0.47)

Stock (post-training). We utilize the individually fine-tuned weights as we conducted in §2, using the same training settings with different random seeds for each model. We report the performance and distance from the pseudo-center of Model Stock (post-training) in Table 7. On the left side of the table, we provide the performance of its counterpart, a uniform averaging of N models ($\mathbf{w}_{\text{avg}}^N$). The improvements from the uniform averaging to Model Stock (post-training) are denoted in the table’s parentheses. The results show that Model Stock (post-training) archives improved distribution shift scores with closer distances toward the center than its counterpart.

5 Related Work

We discuss related works and highlight how our method differs and contributes to the existing works.

Model Soups [35] is a straightforward weight averaging method that merges weights from various fine-tuned models trained with different hyper-parameters. It demonstrates improved in-distribution (ID) and out-of-distribution (OOD) performance. While effective, model soups typically require a large number of fine-tuned models. Our method aims to achieve similar or superior performance improvements more efficiently, utilizing significantly fewer fine-tuning costs. We provide further discussion about Model Soups with our new interpretation in the Appendix.

Robust Fine-tuning. When fine-tuning generalist models like CLIP, we often observe the fine-tuned models lose the generalization ability of the original CLIP model, with decreased OOD performance. To address this issue, several robust fine-tuning approaches have been proposed. LP-FT [16] attempts to preserve pre-trained weights by initially training only a linear probing layer. WiSE-FT [36] improves OOD performance through linear interpolation between fine-tuned and pre-trained weights. While our method shares similarities with WiSE-FT in using pre-trained weights, our method determines interpolation ratios layer-wise based on geometric properties. Focusing on OOD performance, methods suggesting

improved training objectives [7, 23–25, 32, 33] has been proposed. Our approach differs from these methods as we do not propose a new fine-tuning loss. Instead, we perform two fine-tunings and achieve robust performance by merging them.

Weight Center and Flat minima Recent machine learning research has extensively explored the significance of finding flat minima for improved generalization [3, 13, 18, 22]. Keskar *et al.* [14] and Hochreiter & Schmidhuber [11] demonstrated that sharp optima by large batch SGD have steep, harmful directions, while broader optima enhance generalization. Stochastic Weight Averaging (SWA) [13] targets the center of flat minima, enhancing robustness against shifts in the loss landscape between training and test datasets. SWAG [22] builds on SWA by incorporating Bayesian model averaging with a Gaussian posterior to further boost performance. SWAD [3] found that the generalization gap between flat and sharp minima is more pronounced in OOD scenarios than in ID ones. Our method theoretically extends these approaches by efficiently identifying the center of flat minima with novel geometric properties, leading to significant improvements in both ID and OOD performances.

Model Weight Merging. Recent research has explored merging models fine-tuned on various tasks. Methods such as Task Arithmetic [12] and TIES [37] have been proposed. They are also based on the difference between fine-tuned and pre-trained weights (often referred to as the “task vector”). However, our method distinguishes itself through geometric analysis for weight merging. In the domain of Large Language Models (LLMs), merging techniques like WARM [27] and WARP [28] have emerged. Our method has the potential for extension to these areas, offering new avenues for future research.

6 Conclusion

Our study illuminated the fine-tuning process in machine learning, revealing that fine-tuned models’ weights generally exhibit the properties of a Gaussian distribution. The proximity of these models to the center of weights was crucial for improved performance in target domains like ImageNet and under diverse distribution shifts. Utilizing a pre-trained model as a robust anchor point, we efficiently minimized the variance with fewer fine-tuned models, eliminating the need for additional training to find the optimal interpolation ratio.

Additionally, our findings suggested further knowledge and applicability to the models near flat minima and will offer new insights on model weight merging methods. As the pretraining-finetuning paradigm gains more prominence, our insights will provide a foundation for better understanding and optimizing the fine-tuning process in both academia and industry.

Limitation. Due to resource limitations, we could not conduct larger-scale models such as ViT-G. Exploring this will be part of our future work.

References

1. Barbu, A., Mayo, D., Alverio, J., Luo, W., Wang, C., Gutfreund, D., Tenenbaum, J., Katz, B.: Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *NeurIPS* (2019)
2. Beyer, L., Hénaff, O.J., Kolesnikov, A., Zhai, X., Oord, A.v.d.: Are we done with imagenet? *arXiv preprint arXiv:2006.07159* (2020)
3. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: Swad: Domain generalization by seeking flat minima. *NeurIPS* **34**, 22405–22418 (2021)
4. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *CVPRW*. pp. 702–703 (2020)
5. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2021)
6. Gouk, H., Hospedales, T.M., Pontil, M.: Distance-based regularisation of deep networks for fine-tuning. *arXiv preprint arXiv:2002.08253* (2020)
7. Goyal, S., Kumar, A., Garg, S., Kolter, Z., Raghunathan, A.: Finetune like you pretrain: Improved finetuning of zero-shot vision models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19338–19347 (2023)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR*. pp. 770–778 (2016)
9. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: *ICCV*. pp. 8340–8349 (2021)
10. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. In: *CVPR*. pp. 15262–15271 (2021)
11. Hochreiter, S., Schmidhuber, J.: Flat minima. *Neural computation* **9**(1), 1–42 (1997)
12. Ilharco, G., Ribeiro, M.T., Wortsman, M., Schmidt, L., Hajishirzi, H., Farhadi, A.: Editing models with task arithmetic. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=6t0Kwf8-jrj>
13. Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., Wilson, A.G.: Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407* (2018)
14. Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On large-batch training for deep learning: Generalization gap and sharp minima. In: *International Conference on Learning Representations* (2017), <https://openreview.net/forum?id=H1oyRlYgg>
15. Krizhevsky, A.: Learning multiple layers of features from tiny images. In: *Tech Report* (2009)
16. Kumar, A., Raghunathan, A., Jones, R., Ma, T., Liang, P.: Fine-tuning can distort pretrained features and underperform out-of-distribution. *arXiv preprint arXiv:2202.10054* (2022)
17. Ledoux, M.: The concentration of measure phenomenon. No. 89, *American Mathematical Soc.* (2001)
18. Li, H., Xu, Z., Taylor, G., Studer, C., Goldstein, T.: Visualizing the loss landscape of neural nets. *NeurIPS* **31** (2018)

19. Li, T., Huang, Z., Tao, Q., Wu, Y., Huang, X.: Trainable weight averaging: Efficient training by optimizing historical solutions. In: ICLR (2022)
20. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: CVPR. pp. 11976–11986 (2022)
21. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
22. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for bayesian uncertainty in deep learning. *NeurIPS* **32** (2019)
23. Mao, X., Chen, Y., Jia, X., Zhang, R., Xue, H., Li, Z.: Context-aware robust fine-tuning. *IJCV* (12 2023). <https://doi.org/10.1007/s11263-023-01951-2>
24. Nam, G., Heo, B., Lee, J.: Lipsum-ft: Robust fine-tuning of zero-shot models using random text guidance. In: The Twelfth International Conference on Learning Representations (2024)
25. Oh, C., Lim, H., Kim, M., Han, D., Yun, S., Choo, J., Hauptmann, A., Cheng, Z.Q., Song, K.: Towards calibrated robust fine-tuning of vision-language models (2024), <https://arxiv.org/abs/2311.01723>
26. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML (2021)
27. Rame, A., Vieillard, N., Hussenot, L., Dadashi, R., Cideron, G., Bachem, O., Ferret, J.: WARM: On the benefits of weight averaged reward models. In: Forty-first International Conference on Machine Learning (2024), <https://openreview.net/forum?id=s7RDnNUJy6>
28. Ramé, A., Ferret, J., Vieillard, N., Dadashi, R., Hussenot, L., Cedoz, P.L., Sessa, P.G., Girgin, S., Douillard, A., Bachem, O.: Warp: On the benefits of weight averaged rewarded policies (2024), <https://arxiv.org/abs/2406.16768>
29. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imagenet classifiers generalize to imagenet? In: ICML (2019)
30. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *IJCV* **115**(3), 211–252 (2015)
31. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR. pp. 1–9 (2015)
32. Tian, J., He, Z., Dai, X., Ma, C.Y., Liu, Y.C., Kira, Z.: Trainable projected gradient method for robust fine-tuning. In: CVPR. pp. 7836–7845 (2023)
33. Tian, J., Liu, Y.C., Smith, J.S., Kira, Z.: Fast trainable projection for robust fine-tuning. In: *NeurIPS* (2023)
34. Wang, H., Ge, S., Lipton, Z., Xing, E.P.: Learning robust global representations by penalizing local predictive power. In: *NeurIPS* (2019)
35. Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al.: Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: ICML. pp. 23965–23998. PMLR (2022)
36. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., et al.: Robust fine-tuning of zero-shot models. In: CVPR. pp. 7959–7971 (2022)
37. Yadav, P., Tam, D., Choshen, L., Raffel, C., Bansal, M.: TIES-merging: Resolving interference when merging models. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=xtaX3WyCj1>

38. Zaken, E.B., Ravfogel, S., Goldberg, Y.: Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. arXiv preprint arXiv:2106.10199 (2021)