

ERM++: An Improved Baseline for Domain Generalization

Piotr Teterwak[‡] Kuniaki Saito[‡] Theodoros Tsiligkaridis[†] Kate Saenko[‡] Bryan A. Plummer[‡]
 Boston University[‡] MIT Lincoln Laboratory[†]
 {piotr^t, keisaito, saenko, bplum}@bu.edu ttsili@mit.edu

Abstract

Domain Generalization (DG) aims to develop classifiers that can generalize to new, unseen data distributions, a critical capability when collecting new domain-specific data is impractical. A common DG baseline minimizes the empirical risk on the source domains. Recent studies have shown that this approach, known as Empirical Risk Minimization (ERM), can outperform most more complex DG methods when properly tuned. However, these studies have primarily focused on a narrow set of hyperparameters, neglecting other factors that can enhance robustness and prevent overfitting and catastrophic forgetting, properties which are critical for strong DG performance. In our investigation of training data utilization (i.e., duration and setting validation splits), initialization, and additional regularizers, we find that tuning these previously overlooked factors significantly improves model generalization across diverse datasets without adding much complexity. We call this improved, yet simple baseline ERM++. Despite its ease of implementation, ERM++ improves DG performance by over 5% compared to prior ERM baselines on a standard benchmark of 5 datasets with a ResNet-50 and over 15% with a ViT-B/16. It also outperforms all state-of-the-art methods on DomainBed datasets with both architectures. Importantly, ERM++ is easy to integrate into existing frameworks like DomainBed, making it a practical and powerful tool for researchers and practitioners. Overall, ERM++ challenges the need for more complex DG methods by providing a stronger, more reliable baseline that maintains simplicity and ease of use. Code is available at https://github.com/piotr-teterwak/erm_plusplus

1. Introduction

Domain Generalization (DG) [6, 39] tackles the challenge of developing models that can generalize to previously unseen data distributions without relying on the availability of new data distributions for model updates. This is vital when gathering new domain-specific data is impractical and differences between training and deployment

Critical to not catastrophically forget pre-training




Figure 1. Our goal is to provide a simple yet effective baseline that enhances robustness without adding complexity. Pre-training data (green) is often more similar to target data (purple) than source data (blue). Preventing catastrophic forgetting and overfitting to the source is thus critical in DG. We introduce ERM++, which addresses this through three key principles: Training Data Utilization (Sec. 3.1), Initialization (Sec. 3.2), and Regularization (Sec. 3.3).

data are unknown. In multi-source DG, each training sample is categorized as belonging to one of several domains. Many advanced methods extend Empirical Risk Minimization (ERM, [52]) by using this domain membership information [19, 37, 61, 64].

However, recently, DomainBed [21] conducted a comprehensive evaluation of these methods and revealed that ERM with hyper-parameter tuning outperformed the best numbers reported in the literature at the time. This shows the critical importance of well-tuned baselines; they ensure meaningful research results. Nevertheless, ERM [21] tunes a small subset of important factors (Figure 2) and overlooks others which prevent overfitting and catastrophic forgetting.

Yet, the risk of overfitting and catastrophic forgetting remains a significant challenge in DG. DG methods divide training into two phases: pre-training and fine-tuning. Typically, DG methods take an off-the-shelf model pre-trained on some large-scale and generic dataset and fine-tune it on source data. Overfitting to source data is an especially large risk in DG; pre-training data distributions may be more sim-

We tune ERM++ to not overfit or catastrophically forget		
	ERM [21]	ERM++ (ours)
Drop-Out	✓	✓ From [21]
Learning Rate	✓	✓ From [21]
Batch Size	✓	✓ From [21]
Weight Decay	✓	✓ From [21]
Training Data Utilization	✗	✓
Initialization	✗	✓
Regularization	✗	✓

Figure 2. Prior baselines like ERM [21] tune a small set of hyper-parameters. We extend tuning to other factors that control catastrophic forgetting and overfitting to the source.

ilar to evaluation data than the available data for fine-tuning (See Figure 1). For example, a DG model may be pre-trained on photographs, fine-tuned on sketches, and then evaluated on photographs again. Then, the model can easily forget pre-training data when fine-tuned on sketches, which is clearly detrimental to performance. Tuning settings that control overfitting to the source is the focus of our improved baseline, which we call ERM++. In extensive experiments, we show that this improves not only baseline performance over ERM [21] but also outperforms SOTA methods. We organize our additional tuning around three main principles: Training Data Utilization (Section 3.1), Initialization (Section 3.2), and Regularization (Section 3.3).

Training Data Utilization. Modifying the number of samples seen during fine-tuning impacts overfitting. Currently, the DG standard is to cap the number of training steps for each method at a relatively low number while using a validation accuracy to select an exact checkpoint [8, 9, 21]. While this can be effective in preventing strong overfitting, some methods may simply not converge. In addition, it can also limit the effectiveness of other regularizers by shortening training trajectories (see discussion of Model Parameter Average (MPA) Section 3.3). Instead, we tune other hyper-parameters to prevent overfitting while automatically determining the training length and learning rate schedule customized for each model to help ensure convergence. We also use the training data more effectively than the DG standard. Increasing data size is known to improve generalization. Instead of splitting off a validation set and never using it to update the model, we do two passes. The first is to select the number of training steps with a validation set, and the second is to retrain for the selected number of steps with joined training and validation sets, increasing effective dataset size.

Initialization. We consider how to initialize network parameters. Good initialization is known to impact convergence time [41], and strong pre-training procedures can minimize the distance from initialization to strong param-

eters for OOD domains. While it has been shown to be critical for domain adaptation and generalization [1, 29], initialization selection has not yet become part of a standard DG benchmark. We also investigate the role of pre-training data similarity to evaluation (target) data and find that pre-training data similarity is important for DG performance. Still, good initialization can help performance even for datasets very different from pre-training (Section 5).

Regularization. We expand prior works [21] exploration of regularization. Selecting which layers to tune and in which order is an important consideration in how far a fine-tuned network ends up from its initialization. A common example is first training the linear probe [28, 60], but this has not been integrated into a commonly used multi-source DG baseline. Another form of weight-space regularization is averaging model iterates. This results in converging to flatter minima and improved generalization [2, 8, 26, 56]. We investigate different averaging strategies in Section 4.3.

Put together, *we find that the additional tuning in our improved baseline ERM++ outperforms both the prior ERM baselines and all recent SOTA methods on DomainBed*. ERM++ composes well with recent methods and has an attractively low compute cost compared to alternatives (Section 4.4). Applied to ViT-B models, a setting [21] did not explore, ERM++ outperforms ERM by 15%. ERM++ even matches the performance of methods that leverage large-scale text encoders (CAR-FT [38]) and update model parameters with test data (UniDG [62]). Overall, our work underscores the importance of revisiting fundamental methods of controlling overfitting and challenges the need for more complex DG methods.

2. Related Work

Generalization in Neural Networks: Generalization for neural networks has a long research history. Methods include ensembling [17, 25], scaling of data [31], initialization for transfer learning [14, 31, 42], data augmentation [12], and regularizing towards the initialization [58]. Researchers have tried to correlate measures on the training with generalization to the test set [27]. While many of these studies have been for in-domain generalization, their key ideas remain applicable for out-of-domain generalization.

Expanding Training Data: Modifying the amount of training is a critical knob which controls generalization. Data augmentation is a common tool to expand the training data [23, 59, 63, 64]. For example, Inter-domain mixup [59] blends the images of different domains, and augmentation with style transfer can further diversify training images [63]. We leverage a pre-trained model which already was trained with aggressive augmentation, AugMix [23], but integrating stronger augmentation into ERM++ training is a promising direction for future work.

Preventing catastrophic forgetting of good initializations with regularizers:

Several recent approaches aim to leverage generalizable features from a model pre-trained on large-scale data. Adapting such a model to the downstream task without forgetting its generalizable representations is the key to achieving generalization [2, 8, 56] keep a running average of model parameters during training, which can be treated as an ensemble of the initialized model and fine-tuned model. [32] and [60] mitigate feature distortion by pre-training a linear probe first before fine-tuning the backbone, warmstarting the fine-tuning with a good initialization. MIRO [9] maximizes the mutual information in feature space between the fine-tuned and pre-trained networks. We treat warm start and parameter averaging as hyper-parameters to add to ERM++ since their ease of implementation makes them suitable for a baseline.

Domain-invariant feature learning: In multi-source domain generalization, it is common to leverage the domain labels to learn domain-invariant features. CORAL [50] aligns second-order statistics of different domains. DANN [19] uses an adversarial loss to match feature distributions across source domains. However, these methods were evaluated with inconsistent baselines, and when ERM [21] accounted for this, advanced methods such as the above no longer showed a consistent advantage. We develop an even stronger baseline called ERM++.

3. ERM++: Improved ERM baseline for domain generalization

We study the problem of Multi-Source DG for classification. We fine-tune a pre-trained model on data consisting of multiple domains and evaluate it on data from unseen domains. More formally, let us consider training domains $d \in \{d_1, \dots, d_n\}$. Each domain corresponds to a data source that is different from other domains; for example, domains may be "sketch" versus "real" images or different camera locations for wildlife cameras. A training dataset is constructed using all sample, label pairs in all training domains $D = \{(X^{d_1}, Y^{d_1}) \dots (X^{d_n}, Y^{d_n})\}$. After initializing classifier f with a pretrained model, and finetuning on D , it is tested on a held-out testing domain d_{test} . We perform simple empirical risk minimization (ERM), formalized as minimizing the average loss over all samples $\frac{1}{n} \sum_{i \in D} \ell(x_i, y_i)$. We use mini-batch gradient descent and compose batches to be equal parts of each source domain.

Our goal is to investigate the general training components that go into creating an ERM model to provide a strong baseline for future work, ensuring that improvements reported by new methodologies cannot be achieved using simpler means. These include training data utilization (Section 3.1), initialization (Section 3.2), and regularization methods that help prevent overfitting to the source domains (Section 3.3). We call our stronger baseline ERM++.

3.1. Training Data Utilization: Setting Training Length Hyperparameters

A key component of training any neural network is utilizing the (often limited) training data effectively so that both underfitting and overfitting are avoided. A common practice in the domain generalization literature is to split source datasets into (often 80%/20%) train/validation sets under a fixed number of iterations for each dataset (e.g., [2, 8, 21, 45]). The validation data is used to set hyperparameters and perform checkpoint (no. training steps) selection to strike a balance between underfitting and overfitting. This approach has two major drawbacks. First, by creating a separate validation set, a significant portion of labeled data is sacrificed, and data quantity is known to be important for generalization. Second, methods are trained under a fixed (relatively small) cap on number of iterations, resulting in underfitting. To address these issues, we use a two stage training pipeline, where we set training length hyperparameters in the first stage using a validation set, including training length and creating a learning rate schedule. Then, in the second stage, we retrain the network using the entire dataset. Additional details follow.

Automatically determining training length (Auto-LR):

We observe that source validation performance does not saturate on many datasets using the settings in evaluation benchmarks used in prior work (See Supplementary), resulting in under-trained models. We avoid this issue by automating the selection of training length as well as the learning rate schedule used on each dataset. Specifically, we begin by splitting a dataset into its training and validation sets, which we use to set training length hyperparameters. To determine the training length, we keep training a model until it does not improve performance on a validation set (held out from source domains) within the previous N steps. We then reduce the learning rate by 1/10th. Learning rates are reduced on the validation performance plateau twice, and we stop on the third drop (see supplementary for pseudocode). As this can provide variable training times, we do provide a high training step cap as an additional stopping condition (4x larger than DomainBed). Stopping after the third decrease helps avoid overfitting on some datasets while also allowing for longer training times and also provides a longer optimization trajectory for **Model Parameter Averaging** (See Section 3.3).

Using the full data (Full Data, FD): After setting our training length and learning rate schedule in the first stage, stage two retrains the model using all our learned hyperparameters with the entire (train+validation) dataset. This maximizes the amount of training data we use for a model to ensure we report its best performance with the available data. In our experiments, we only set the training length and learning rate schedule once per model and reused those parameters for any other required tuning.

3.2. Initializing Model Weights

Domain generalization methods do not train a model from scratch, but rather transfer the weights of an existing model, *e.g.*, ImageNet [13] pretraining. However, prior work assumes TorchVision model weights (*e.g.*, ERM [21]) for ResNet-50, which were trained using an outdated and basic training recipe. For ViT-B models, initializations beyond OpenAI CLIP [44] are rarely considered. Thus, another question we explore is what effect the data or training strategy has on DG performance. We refer to the best initialization method from our analysis as **Strong Init**, which represents the Augmix [23] weights for ResNet-50 and DinoV2 [42] weights for ViT-B.

3.3. Regularization

Regularization has long been used to prevent the overfitting of models to training data. Overfitting is especially problematic in DG since the source data has a different distribution than the target distribution. In fact, the pre-training data may be more similar to the target data than the source data. Thus, we explore several regularization techniques and focus on those which do not require an additional hyperparameter λ in the loss.

Model Parameter Averaging (MPA) averages model iterates [2, 8, 26, 36, 45, 46, 55, 56], which improves generalization by converging to flatter minima [26]. Arpit *et al.* [2] use a simple method for parameter averaging where simply all iterates are averaged (MPA). Not only does MPA converge to flatter minima, but it also averages model parameters from near the initialization and from near the converged minima, acting as regularizer by pulling the model closer to initialization. Furthermore, averaged parameters can be viewed as an ensemble [25]. We verify that MPA works in combination with other techniques present in ERM++. Finally, long training enabled by **Auto-lr** (another ERM++ component, Section 3.1) allows the optimization path to lengthen. This can improve the effectiveness of averaging model weights by increasing the diversity of averaged model parameters [25]. This is particularly intuitive with an ensemble view of weight averaging, as diverse ensembles are more performant than homogenous ones.

Initializing Classifier Weights (Warm Start, WS): New class labels require a new classification layer, typically via random initialization. However, a recurring observation made by many researchers is that a model may suffer from divergence from the initialization due to the noisy, and large magnitude, gradients from newly initialized layers [20, 22, 45]. In the case of pretrained models, this results in catastrophic forgetting of robust, pre-trained features. To address this, Warmstart is utilized (WS) [28, 60] (also commonly referred to as warmup), where the new layer weights are trained with all pretrained weights kept frozen for several hundred steps. After this short training cycle, new and

old layer weights are finetuned together (sometimes except for BatchNorm layers). This acts as a regularizer because after training the linear classifier, gradient norms are much smaller, resulting in parameters that are closer to the initialization after training.

Unfreezing BatchNorm for ResNets (UBN): Batch norm running statistics are frequently initialized from a pretrained model and kept frozen (*i.e.* using batch normalization in evaluation mode) [8, 9, 21]. However, using training batch statistics (*i.e.* using batch normalization in train mode) adds noise to the training, and acts as a regularizer. Therefore, we experiment with unfreezing batch norm.

Attention Tuning for ViTs (Attn Tune): Regularization methodologies such as Warm Start and Model Parameter Averaging are still applicable for ViT’s, but Vision Transformers also offer additional options for weight regularization due to their unique architecture and properties. [18] shows that the feed-forward features in Vision Transformers are hardly useful and that classification is done in the attention layers. Supporting this, [51] shows that attention-tuning is effective for standard transfer learning. By tying all weights that are not part of the attention layers explicitly to their initialized values, *Attention Tuning* prevents many model parameters from being finetuned and, therefore, acts as a regularizer. We find it to be effective for ViTs.

4. Experiments

We evaluate ERM++ using top-1 classification accuracy on a diverse set of datasets commonly used in multi-source DG, including **OfficeHome** [53], **DomainNet** [43], **PACS** [34], **VLCS** [16], and **TerraIncognita** [5]; data and implementation are detailed in the Supplementary.

Implementation details. We follow the DomainBed training procedure, which is used for the ERM baseline, and add additional components for ERM++. In particular, we use the default hyper-parameters from DomainBed [21], *e.g.*, batch size of 32 (per-domain), learning rate of $5e-5$, ResNet dropout value of 0, and weight decay of 0. We train on all source domains except one, validate the model on held-out data from the sources every 300 steps, and evaluate on the held-out domain. The validation set does not contain any target data to avoid data leakage.

Results. Table 1 reports the performance of DG methods built on ERM across five DomainBed datasets. Comparing the results using ERM++ on a ResNet backbone in Table 1(b) to methods from prior work in Table 1(a), we see ERM++ boosts performance by nearly 1%. Within Table 1(b) we see this boost can be increased by another 0.5% by combining ERM++ with methods from prior work. However, this gain is much smaller than methods like MIRO or DIWA combined with ERM (around a 4 point gain as seen in Table 1(a)). We also note that MIRO and DIWA are more computationally expensive than ERM++ during

	OfficeHome	PACS	DomainNet	Incognita	VLCS	Avg.
(a) ResNet-50 Backbone; built on ERM [21]						
MLDG [33]	66.8±0.6	84.9±1.0	41.2±0.1	47.7±0.9	77.2±0.4	63.6
ERM [52]	67.6±0.2	84.2±0.1	44.0±0.1	47.8±0.6	77.3±0.1	64.2
ERM+ MMD [35]	66.3±0.1	84.7±0.5	23.4±9.5	42.2±1.6	77.5±0.9	58.8
ERM + DANN [19]	65.9±0.6	83.6±0.4	38.3±0.1	46.7±0.5	78.6±0.4	62.6
ERM + Mixup [57, 59]	68.1±0.3	84.6±0.6	39.2±0.1	47.9±0.8	77.4±0.6	63.4
ERM + CORAL [50]	68.7±0.3	86.2±0.3	41.5±0.1	47.6±1.0	78.8±0.6	64.5
SagNet [40]	68.1±0.1	86.3±0.2	40.3±0.1	48.6±1.0	77.8±0.5	64.2
mDSDI [7]	69.2±0.4	86.2±0.2	42.8±0.1	48.1±1.4	79.0±0.3	65.1
ERM + MIRO [9]	70.5±0.4	85.4±0.4	44.3±0.2	50.4±1.1	79.0±0.0	65.9
ERM + SWAD [8]	70.6±0.2	88.1±0.1	46.5±0.1	50.0±0.3	79.1±0.1	66.9
ERM + CORAL + SWAD [50]	71.3±0.1	88.3±0.1	46.8±0.0	51.0±0.1	78.9±0.1	67.3
ERM + DIWA [45]	72.8	89.0	47.7	51.9	78.6	68.0
ERM + MIRO + SWAD [9]	72.4±0.1	88.4±0.1	47.0±0.0	52.9±0.2	79.6±0.2	68.1
(b) ResNet-50 Backbone; built on ERM++						
ERM++ (w/out <i>FD</i>)	74.0	89.6	45.8	51.0	78.3	67.7
SWAD + ERM++ (w/out <i>FD</i>)	74.1	89.6	49.7	49.7	77.3	68.0
ERM++	74.7±0.0	89.8±0.3	50.8±0.0	51.2±0.3	78.0±0.1	68.9
DIWA + ERM++	75.1	90.0	51.5	51.4	78.6	69.3
MIRO + ERM++	76.3	88.8	50.4	53.4	77.9	69.4
(c) ViT-B Backbone						
ERM [52]	66.4	83.4	44.4	35.3	75.9	61.1
ERM + MIRO [9]	82.5	95.6	54.0	54.3	82.5	73.7
ERM + CAR-FT [38]*	85.7	96.8	62.5	61.9	85.5	78.5
ERM + UNI-DG [62]†	86.2	96.7	61.3	62.4	86.3	78.6
ERM++	86.7	96.8	59.8	67.4	82.4	78.6

Table 1. **ERM++ Domain Generalization Performance:** (a) We show DG performance of methods with an ERM foundation. [21] provided an ERM baseline, on which new methods like MIRO and SWAD could be built. All methods below the horizontal line were developed after the tuned ERM baseline [21] was published, showing how it pushed the DG field forward. ERM++ provides a refreshed baseline for future work to do the same. (b) Performance of our ERM++ baseline. We show that ERM++ is an even stronger foundation for recent methods than ERM. SWAD is built on top of ERM++ w/out FD because it requires validation data for interval selection. (c) When we extend ERM++ to vision transformers, ERM++ outperforms ERM by a full 15%, showing the generality of ERM++ principles. *CAR-FT [38] utilizes the text encoder of CLIP for regularization, which adds 63 million parameters to the model and is not present in all ViT-B models. †UNI-DG utilizes target data for model updates, so it falls outside of the bounds of DG, but we add it for reference.

training (Section 4.4), and therefore ERM++ alone may be preferable in compute-constrained settings.

Comparing the methods using a ViT-B backbone in Table 1(c), we see that ERM++ can match performance with OpenAI CLIP initialized methods that have an unfair advantage by either increasing their model size via a text encoder (CAR-FT), or use target data for model updates (UNI-DG). Under a fairer comparison, ERM++ with a ViT-B backbone outperforms prior work by 5-17%. This demonstrates that by optimizing the training pipeline in ERM++, we can greatly boost performance, matching or exceeding the accuracy of more complicated, state-of-the-art methods.

Below we analyze ERM++ components: training length (Section 4.1), initialization (Section 4.2), and regularization (Section 4.3). Section 4.4 compares ERM++’s computational cost to prior work. Section 5 evaluates the effect of pretraining data and how well methods generalize to far out-of-domain settings using non-web-scraped data.

4.1. Training Data Utilization

Automatic Training Length/Learning Rate Schedule (*Auto-LR*): Section 3.1 describes our two stage training pipeline. In this first stage, we set the training length and learning rate schedule automatically using a validation set for each method. This helps ensure we train until convergence, which can vary for each method and dataset. Training until convergence can have a large impact on transfer learning performance [10], and model averaging performs best when weights are diverse [45], which benefits from the additional training time. At the same time, automatic stopping prevents over-fitting. In Table 2(a) we show this boosts performance by 0.5 on average.

Using the full data (*FD*): The most common ERM implementation [21] splits the training data into 80/20 train/validation splits, significantly reducing the size of the training set. Thus, our second stage of training recombines the train/validation split and re-trains the model using the

Backbone/Category	Ablation	OfficeHome	PACS	DomainNet	TerraIncognita	VLCS	Avg.
(a) ResNet-50 Backbone	ERM++	74.7	89.8	50.8	51.2	78.0	68.9
Train Length	w/o <i>Auto-LR</i>	74.6	87.9	49.8	51.1	78.6	68.4
Train Length	w/o <i>FD</i>	74.0	89.6	45.8	51.0	78.3	67.7
Initialization	w/o <i>S. Init</i>	72.6	88.8	48.6	49.2	78.7	67.6
Regularization	w/o <i>UBN</i>	74.7	90.1	49.9	49.0	78.6	68.3
Regularization	w/o <i>WS</i>	74.3	89.5	50.7	51.9	77.2	68.7
Regularization	w/o <i>MPA</i>	71.5	72.3	46.9	48.9	79.8	65.9
(b) ViT B/16 Backbone	ERM++	86.7	96.8	59.8	67.4	82.4	78.6
Regularization	w/o <i>Attn Tune</i>	85.6	95.6	60.3	62.6	81.1	77.0

Table 2. We present the overall ablation for ERM++. (a) Each component contributes to the strong performance of ERM++ with *Model Parameter Averaging* being the most impactful part of ERM++. The use of *Full Data*, *Model Parameter Averaging*, and *Strong Init* are most critical, while the other methods give small boosts which add up (see cumulative study in the supplementary). (b) We also ablate the ViT-specific component of *Attention Tuning*. We find that *Attention Tuning* improves performance on almost datasets, except for DomainNet, which has the largest amount of data. The jump is especially large on TerraIncognita, of about 5%.

hyperparameters set in the first stage. In Table 2(a) we show that this improves performance by up to 5% on DomainNet. We emphasize that this is a fair comparison to prior work; validation data used for hyperparameter tuning (as in prior work) should also be considered training information.

4.2. Initializing Model Weights

In Section 3.2 we note that prior work has often initialized their model weights using old models trained with out-of-date techniques. In this section, we explore the effect of using **Stronger initializations** (*S. Init*) of model weights using more recent methods. Table 2(a) shows can obtain a 1% boost on average, although we note that ERM++ still matches the performance of prior work reported in Table 1(a) when using the same initialization, but without requiring adding new loss functions requiring hyperparameter tuning (as in MIRO [9]) or complicated training procedures (as with DIWA [45]). However, selecting the best model takes some care. To illustrate, let us consider the following initializations for ResNet-50. **TorchVision Model Weights** is the standard ImageNet pretrained initialization present in TorchVision. It was trained with weak augmentations for 90 epochs. **AugMix** [23] is a method used to improve model consistency using augmentations without training the model on data which is too different from the test data. AugMix takes two augmented views of an image and mixes them in pixel space. Then, the model is trained to produce consistent output between two AugMix augmentations and the clean image. **ResNet A1** initializes weights from the training recipe presented in [54]. The model is heavily tuned to find training settings that result in very strong ImageNet performance. Examples include training for 600 epochs, the LAMB optimizer, strong augmentations, and a binary cross-entropy. **MealV2** [48] is a highly performant ensemble,

distilled into a ResNet-50. In particular, a SeNet-154 [24] (81.23% ImageNet Top-1) and ResNet-152 (81.02% ImageNet Top-1) are distilled into ResNet-50.

For ViT-B backbones, we consider the following. **ImageNet 1k** [49]: A recipe which is found to provide good results on ImageNet, making the model directly comparable to ImageNet pre-trained ResNet-50 model in terms of pre-training data similarity to Domain Generalization targets. **ImageNet-21k** [15] is the original ViT model. While the data scale is larger than ImageNet-1k, the pre-training domain is still real images. **OpenAI CLIP ViT-B/16** [44] is a strong ViT-B-16 model contrastively trained on 400 million image-caption pairs scraped from the web. **DinoV2 ViT-B/14** [42] is a self-supervised training method, trained on a highly curated dataset of 142 million samples.

Table 3 reports the performance of each initialization method, where we find that selecting a good checkpoint is not as simple as benchmarking ImageNet performance. In particular, both the ResNet-50 A1 and Meal V2 weights achieve much better ImageNet Top-1 Accuracy than the standard TorchVision weights but do not achieve the best DG performance. The DG performance of the AugMix weights makes it a reasonable choice.

We also observe that model distillation, which strongly improves source accuracy, does not increase overall DG performance. Meal V2 is a distillation of the ensemble of two very strong ImageNet models into a ResNet-50. Interestingly, the student in Meal V2 is initialized with the same AugMix-trained network as our experiments. Thus, the differences in performance can be strictly attributed to the effects of model distillation. Detailed results in the Supplementary show that ImageNet-like domains improve while performance on other domains suffers. This suggests that the distillation process effectively matches the student to

	Avg	ImageNet
(a) ResNet-50 Backbone		
TorchVision Weights	66.6	76.1
AugMix [23]	68.4	79.0
Meal V2 [48]	68.3	80.7
ResNet A1 [54]	62.3	80.4
(b) ViT-B Backbone		
ImageNet [15]	65.3	79.2
ImageNet-22k [49]	67.5	85.1
OpenAI CLIP [44]	76.9	80.2
DINOv2 [42]	78.6	84.5

Table 3. **DG performance with different ResNet-50 and ViT-B initialization.** (a) **ResNet-50:** The differences between initializations are very substantial. Interestingly, improved ImageNet accuracy does not strongly correlate with improved performance. (b) **ViT-B:** We find initialization is even more important here. ImageNet ViT models are weaker than ImageNet ResNet models, but scaling pre-training data (CLIP and DINO) improves performance by over 10%.

the teacher over the data used in the distillation process, but not elsewhere. AugMix is a model trained with generalization to synthetic corruptions as a goal, and it results in very strong DG performance. Thus, while ImageNet accuracy is not a good indicator of DG performance, investigating the correlation between synthetic corruption performance and DG performance is promising.

From our transformer results in Table 3(b), we generally find scaling pre-training data improves performance. OpenAI CLIP (400 million pre-training data points) outperform ImageNet-21k pretraining (14 million data points) by 10 percent, which improves over ImageNet pretraining (1.4 million data points) by 2 percent. However, training methodology can make up for dataset scale, as DinoV2 was only trained on 142 million data points yet achieves higher DG performance. Still, the large pretraining datasets used by these models can be a large driver of performance by increasing similarity to DG performance(Section 5).

4.3. Regularization

Table 2 reports the effect of the different regularization components we discussed in Section 3.3. We find that each of the regularization components of model parameter averaging (MPA), warmstart (WS), unfreezing the batchnorm (UBN), and attention tuning (Attn Tune) all boosts performance by 0.5-3%. The largest gains come from MPA and Attn Tuning, which improves performance by 3% and 1.5%, respectively. Most of the gain using Attn Tune comes from datasets with less available data, such as TerraIncognita’s 5% boost. We find that attention tuning also results in attention maps that attend to salient features, more so than full fine-tuning and especially more than just the pre-trained model (See Supplementary for examples). We also note

	HP Searches	Train FLOPS	Top-1
ERM++ w/fewer steps	2 (for <i>Auto-LR</i>)	1x	68.4%
ERM++	2 (for <i>Auto-LR</i>)	<4x	68.9%
MIRO [9]	4 (for λ)	2x	68.1%
DIWA [45]	60	15x	68.0%

Table 4. Computational Cost: ERM++ achieves high performance without extensive hyper-parameter searches, instead using reasonable default ones. Even without an increased training step cap (Section 3.1), we’re able to achieve SOTA performance. Train FLOPs are relative to ERM++w/out the increased training step cap of *Auto-LR* Results are reported for ResNet-50.

that by only finetuning the transformer’s attention layers, we also reduce the memory required to train the network.

We also explored a setting where instead of averaging model weights, we attempt to include diversity between the models being averaged as this has been shown to boost performance [45]. Following [36], we first train a generalist model on all source domains for 28k steps, then train specialist models for 28k steps (1 model per domain), before averaging parameters. Although averaging specialists improves over ERM by 2%, it underperforms averaging iterates of a generalist by 2%. More detailed results are reported in the supplementary. One possible explanation is that each domain has too little data.

4.4. ERM++ Computational Cost

Table 4 compares the training cost of ERM++ to prior work. We show that DIWA [45] and MIRO [9] come at significant cost due to training a diverse ensemble or hyperparameter tuning for the contribution λ of its new loss function, respectively. In comparison, the only hyperparameters tuned for each are the training length and learning rate schedule in our two-stage process described in Section 3.1. This results in lower computational cost compared with prior work while also obtaining state-of-the-art results.

5. Generalizing Beyond Web-scraped Datasets

We have demonstrated that ERM++ is a highly effective recipe for DG on a suite of common DomainBed benchmarks. However, OfficeHome, PACS, and DomainNet are all web-scraped datasets. Because pre-training of foundation ViTs is also on very large scale web-scraped data, it is important to understand how much of their improved performance is driven by train-test similarity. Therefore, we expand our evaluation to include two out-of-domain (OOD) datasets: FMOW [11, 30] and PCAM [3, 30]. We also measure similarity of LAION-400M [47](a representative pre-training dataset) to DG datasets to provide additional insight. We find that dataset similarity is a strong driver of performance, but that the strong DINOv2 initialization can deliver high performance even on OOD datasets.

	OfficeHome	PACS	DomainNet	TerraIncognita	VLCS	FMOW	PCAM	
(a)	Image Similarity	79.8	77.3	73.6	66.1	69.9	70.4	76.9
	Text Similarity	83.1	85.8	83.2	83.0	86.4	77.2	71.3
	Average Similarity	81.4	81.7	78.3	74.6	78.0	73.7	74.0
	Approximate Duplicates	8%	9%	7.5%	0%	0%	0%	0%
(b)	ERM++ w/ResNet-50 Init.	74.7	89.8	50.8	52.9	79.6	48.1	89.4
	ERM++ w/DINOv2 Init.	86.7	96.8	59.8	67.4	82.4	54.8	95.9
	ERM++ w/OpenAI CLIP Init.	85.9	96.2	61.1	59.5	83.2	48.2	93.7
	OpenAI CLIP zeroshot	82.3	96.1	57.7	23.1	82.4	19.7	55.3

Table 5. **Relationship between dataset similarity and DG performance.** We compute the similarity of LAION-400M for DomainBed (left) and Stanford-WILDS (right) datasets. Overall, performance gains over ImageNet pretraining are largest on datasets where text similarity is high to LAION-400M, suggesting pre-training data plays a key role in DG performance. Nevertheless, DINOv2 init performs strongly on dissimilar data as well. More discussion in Section 5.

Metrics. In addition to hold-one-out DG performance on models with several different initializations, we also measure similarity between target datasets and a representative large-scale web-scraped pre-training datasets, LAION-400M. To do so, we use an approximate nearest neighbors algorithm [4], and build a retrieval index with OpenAI CLIP-B/32 image embeddings. We then measure dataset similarity by querying the index with dataset samples from the DG dataset, and record the cosine similarity of the nearest neighbor. Averaged across 1000 random samples, this constitutes an *image similarity*. We do the same with class labels of the target datasets, to create a *text similarity*. We combine the two modalities by averaging, creating a *average similarity*. We also probe the DG for exact duplicates within LAION-400M, by sampling a random 200 images from each downstream dataset, and retrieving their nearest neighbor using the image index described above. We manually inspect these images for exact duplicates (except for slight crop differences) between query and retrieval.

Results. Table 5(a) reports the similarity between pretraining dataset LAION-400M [47] and the images in the DG benchmark datasets. We observe a number of interesting findings. First, 3 of 4 datasets (DomainNet, OfficeHome, PACS) with the highest average similarity with LAION-data also contain the largest number of duplicates. Second, the average similarities score of DomainBed datasets is much higher than that of the Stanford-WILDS datasets, suggesting that most current benchmarks evaluate Domain Generalization results that are not very different from large-scale pre-training. Third, in Table 5(b) we find that even without fine-tuning with ERM++, CLIP zero-shot performance is very high on all the datasets which have high average similarity. This suggests CLIP is a very strong initialization even without any finetuning. However, ERM++ fine-tuning is critical for all datasets which have poor zero-shot performance. Finally, ERM++ with CLIP and DINOv2 initializations improve most over ImageNet initializations (ResNet-50) on datasets which have high *text similarity* to LAION-400M. The exception is VLCS, which consists of all real domains, and therefore is also similar to ImageNet,

in which case ImageNet pre-training may be sufficiently strong. This suggests that as long as concepts are present in the pre-training, ERM++ fine-tuning with foundational ViTs can overcome image dissimilarity between fine-tuning and pre-training images. Put together, these observations indicate that pre-training data similarity, especially in text-space, to target domains is a big driver of ERM++ performance. Nevertheless, the *Strong Init* (DINOv2) for ViT-B boosts performance over OpenAI clip init on the dissimilar WILDS-FMoW by 7% and WILDS-PCAM by 2%, suggesting that pre-training methodology complements pre-training data composition, and delivers strong performance.

6. Conclusion

This paper develops a strong yet simple baseline, ERM++, for improving the performance of DG models. By identifying several techniques for enhancing ERM, our approach achieves significant gains in DG performance, reporting a 5% boost over ERM and 1% better than the SOTA on ResNet-50. We find ERM++ improves over ERM even more significantly when applied to unstable and over-parameterized ViT models, by over 15%. Our results highlight the importance of improving the training procedure for enhanced DG performance.

Future Extensions of ERM++: While ERM++ demonstrates significant improvements, it also opens up several promising avenues for further exploration. For instance, integrating stronger data augmentation techniques could be a natural extension to improve generalization. Additionally, we leverage hyper-parameters from [21], but searching for better values may further improve performance.

Acknowledgements: DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

This material is based upon work supported by the Under Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Under Secretary of Defense for Research and Engineering.

References

- [1] Simone Angarano, Mauro Martini, Francesco Salvetti, Vittorio Mazza, and Marcello Chiaberge. Back-to-bones: Re-discovering the role of backbones in domain generalization. *arXiv preprint arXiv:2209.01121*, 2022. 2
- [2] Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 2, 3, 4
- [3] Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermsen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 2018. 7
- [4] Romain Beaumont. Clip retrieval: Easily compute clip embeddings and build a clip retrieval system with them. <https://github.com/rom1504/clip-retrieval>, 2022. 8
- [5] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, pages 456–473, 2018. 4
- [6] Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems*, 24, 2011. 1
- [7] Manh-Ha Bui, Toan Tran, Anh Tran, and Dinh Phung. Exploiting domain-specific features to enhance domain generalization. *Advances in Neural Information Processing Systems*, 34:21189–21201, 2021. 5
- [8] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. In *Advances in Neural Information Processing Systems*, volume 34, pages 22405–22418, 2021. 2, 3, 4, 5
- [9] Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pages 440–457. Springer, 2022. 2, 3, 4, 5, 6, 7
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 5
- [11] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 7
- [12] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. 2
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4
- [14] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655. PMLR, 2014. 2
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 6, 7
- [16] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1657–1664, 2013. 4
- [17] Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019. 2
- [18] Yossi Gandelsman, Alexei A Efros, and Jacob Steinhardt. Interpreting clip’s image representation via text-based decomposition. In *The Twelfth International Conference on Learning Representations*, 2023. 4
- [19] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 1, 3, 5
- [20] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 4
- [21] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. 1, 2, 3, 4, 5, 8
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [23] Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. 2, 4, 6, 7
- [24] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018. 6
- [25] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. Snapshot ensembles: Train 1, get m for free. In *International Conference on Learning Representations*, 2016. 2, 4

- [26] Pavel Izmailov, Dmitrii Podoprikin, T. Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *Conference on Uncertainty in Artificial Intelligence*, 2018. 2, 4
- [27] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. *arXiv preprint arXiv:1912.02178*, 2019. 2
- [28] Fahdi Kanavati and Masayuki Tsuneki. Partial transfusion: on the expressive influence of trainable batch norm parameters for transfer learning. In *Medical Imaging with Deep Learning*, pages 338–353. PMLR, 2021. 2, 4
- [29] Donghyun Kim, Kaihong Wang, Stan Sclaroff, and Kate Saenko. A broad study of pre-training for domain generalization and adaptation. In *The European Conference on Computer Vision (ECCV)*, 2022. 2
- [30] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021. 7
- [31] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 491–507. Springer, 2020. 2
- [32] Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pre-trained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022. 3
- [33] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 5
- [34] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 4
- [35] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 5
- [36] Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A Smith, and Luke Zettlemoyer. Branch-train-merge: Embarrassingly parallel training of expert language models. *arXiv preprint arXiv:2208.03306*, 2022. 4, 7
- [37] Ya Li, Mingming Gong, Xinmei Tian, Tongliang Liu, and Dacheng Tao. Domain generalization via conditional invariant representations. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 1
- [38] Xiaofeng Mao, Yufeng Chen, Xiaojun Jia, Rong Zhang, Hui Xue, and Zhao Li. Context-aware robust fine-tuning. *International Journal of Computer Vision*, pages 1–16, 2023. 2, 5
- [39] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International conference on machine learning*, pages 10–18. PMLR, 2013. 1
- [40] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8690–8699, 2021. 5
- [41] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 512–523. Curran Associates, Inc., 2020. 2
- [42] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. 2, 4, 6, 7
- [43] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 4
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 4, 6, 7
- [45] Alexandre Rame, Matthieu Kirchmeyer, Thibaud Rahier, Alain Rakotomamonjy, patrick gallinari, and Matthieu Cord. Diverse weight averaging for out-of-distribution generalization. In *Advances in Neural Information Processing Systems*, 2022. 3, 4, 5, 6, 7
- [46] David Ruppert. Efficient estimations from a slowly convergent robbins-monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988. 4
- [47] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 7, 8
- [48] Zhiqiang Shen and Marios Savvides. Meal v2: Boosting vanilla resnet-50 to 80%+ top-1 accuracy on imagenet without tricks. *arXiv preprint arXiv:2009.08453*, 2020. 6, 7
- [49] Andreas Peter Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *Transactions on Machine Learning Research*, 2022. 6, 7

- [50] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14*, pages 443–450. Springer, 2016. 3, 5
- [51] Hugo Touvron, Matthieu Cord, Alaaeldin El-Nouby, Jakob Verbeek, and Hervé Jégou. Three things everyone should know about vision transformers. In *European Conference on Computer Vision*, pages 497–515. Springer, 2022. 4
- [52] Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999. 1, 5
- [53] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017. 4
- [54] Ross Wightman, Hugo Touvron, and Hervé Jégou. Resnet strikes back: An improved training procedure in timm. *arXiv preprint arXiv:2110.00476*, 2021. 6, 7
- [55] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022. 4
- [56] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7959–7971, 2022. 2, 3, 4
- [57] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 6502–6509, 2020. 5
- [58] LI Xuhong, Yves Grandvalet, and Franck Davoine. Explicit inductive bias for transfer learning with convolutional networks. In *International Conference on Machine Learning*, pages 2825–2834. PMLR, 2018. 2
- [59] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677*, 2020. 2, 5
- [60] Andrew Zhai and Hao-Yu Wu. Classification is a strong baseline for deep metric learning. In *British Machine Vision Conference (BMVC)*, 2019. 2, 3, 4
- [61] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. *Advances in Neural Information Processing Systems*, 34:23664–23678, 2021. 1
- [62] Yiyuan Zhang, Kaixiong Gong, Xiaohan Ding, Kaipeng Zhang, Fangrui Lv, Kurt Keutzer, and Xiangyu Yue. Towards unified and effective domain generalization. *arXiv:2310.10008*, 2023. 2, 5
- [63] Zhun Zhong, Yuyang Zhao, Gim Hee Lee, and Nicu Sebe. Adversarial style augmentation for domain generalized urban-scene segmentation. In *Advances in Neural Information Processing Systems*, 2022. 2
- [64] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *International Conference on Learning Representations*, 2021. 1, 2