

Boosting the Generalization Capability in Cross-Domain Few-shot Learning via Noise-enhanced Supervised Autoencoder

Hanwen Liang^{*1} Qiong Zhang^{*†2} Peng Dai¹ Juwei Lu¹

hanwen.liang@huawei.com, qiong.zhang@stat.ubc.ca, {peng.dai, juwei.lu}@huawei.com

¹Huawei Noah's Ark Lab, Canada. ²Department of Statistics, University of British Columbia, Vancouver, Canada.

Abstract

State of the art (SOTA) few-shot learning (FSL) methods suffer significant performance drop in the presence of domain differences between source and target datasets. The strong discrimination ability on the source dataset does not necessarily translate to high classification accuracy on the target dataset. In this work, we address this cross-domain few-shot learning (CDFSL) problem by boosting the generalization capability of the model. Specifically, we teach the model to capture broader variations of the feature distributions with a novel noise-enhanced supervised autoencoder (NSAE). NSAE trains the model by jointly reconstructing inputs and predicting the labels of inputs as well as their reconstructed pairs. Theoretical analysis based on intra-class correlation (ICC) shows that the feature embeddings learned from NSAE have stronger discrimination and generalization abilities in the target domain. We also take advantage of NSAE structure and propose a two-step fine-tuning procedure that achieves better adaption and improves classification performance in the target domain. Extensive experiments and ablation studies are conducted to demonstrate the effectiveness of the proposed method. Experimental results show that our proposed method consistently outperforms SOTA methods under various conditions.

1. Introduction

After years of development, deep learning methods have achieved remarkable success on visual classification tasks [17, 37, 21, 30]. The outstanding performance, however, heavily relies on large-scale labeled datasets [5]. Meanwhile, although some large-scale public datasets, e.g. ImageNet [9], have made it possible to achieve better than human performance on common objects recognition, practical applications of visual classification systems usually target at categories whose samples are very difficult to collect, e.g.

^{*}Equal contribution with alphabetical order. [†]Work done as an intern in Noah's Ark Lab.

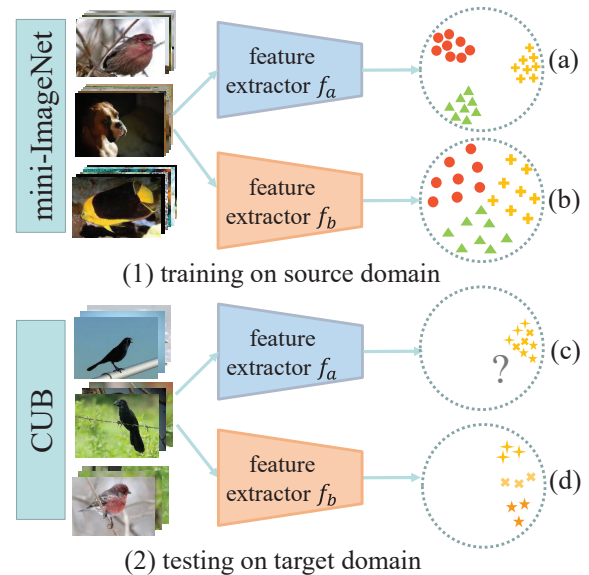


Figure 1: **Motivation illustration.** Visualization of feature embeddings by a less-generalized feature extractor f_a and a well-generalized feature extractor f_b cross source and target domains.

medical images. The scarcity of data limits the generalization of current vision systems. Therefore, it is essential to learn to generalize to novel classes with a limited number of labeled samples available in each class. Cross-domain few-shot learning (CDFSL) is proposed to recognize instances of novel categories in the target domain with few labeled samples. Different from general few-shot learning (FSL) where large-scale source dataset and few-shot novel dataset are from the same domain, target dataset and source dataset under CDFSL setting come from different domains, *i.e.* the marginal distributions of features of images in two domains are quite different [49].

Much work has been done to solve FSL problem and obtained promising results [41, 14, 35, 36, 11, 33, 43]. However, [6, 16] show that the state-of-the-art (SOTA) meta-learning based FSL methods fail to generalize well and perform poorly under CDFSL setting. It is therefore of great

importance to improve the generalization capability of the model and address the domain shift issue from source to target domains. [39] proposes to add a feature-transformation layer to simulate various distributions of image features in training. However, this method requires access to a great amount of data from multiple domains during training. [47] combines the FSL learning objective and the domain adaptation objective, while their basic assumption that source and target domain have identical label sets limits its application. [16] experimentally shows that the traditional transfer learning methods can outperform meta-learning FSL methods by a large margin on the benchmark. In these methods, a feature extractor is pre-trained on the source dataset and then fine-tuned on the target dataset with only a few labeled samples. Following this thread, [26] proposes to regularize the eigenvalues of the image features to avoid negative knowledge transfer.

In this work, our observation is that generalization capability plays a vital role for representation learning in cross-domain settings. As the feature distributions of different domains are distinct, a competent feature extractor on the source domain does not necessarily lead to good performance on the target domain. It may overfit to the source domain and fail to generalize in the target domain. Fig. 1(a) shows an example of a less-generalized feature extractor f_a that fits the source dataset very well and achieves high performance in downstream classification task. When the model is transferred to a different target domain, as shown in Fig. 1(c), the corresponding feature embeddings of different classes may become less discriminative or even inseparable. On the other hand, a less perfect feature extractor f_b on the source domain (Fig. 1(b)), may have stronger generalization capability and obtain more discriminative feature embeddings in the target domain (Fig. 1(d)). Under this intuition, we focus on boosting the generalization capability of the transfer learning based methods, and investigate a multi-task learning scheme that shows the potential to improve generalization performance in [22]. Specifically, we propose a novel noise-enhanced supervised autoencoder (NSAE) that takes more than classification tasks and learns the feature space in discriminative and generative manners. We take advantage of the NSAE structure in the following aspects. First of all, it is shown in [22] that a supervised autoencoder can significantly improve model generalization capability. We develop the model to jointly predict the labels of inputs and reconstruct the inputs. Secondly, motivated by the observation that “the addition of noise to the input data of a neural network during training can, in some circumstances, lead to significant improvements in generalization performance” [31, 2, 1], we consider reconstructed images as noisy inputs and feed them back to the system. The joint classifications based on reconstructed and original images further improve the generalization capability and

avoid the necessity of designing a mechanism to add hand-crafted noises. Thirdly, we develop a two-step fine-tuning procedure to better adapt model to the target domain. Before tuning with the supervised classification method, we first tune model on the target domain in an unsupervised manner by learning to reconstruct images in novel classes. Furthermore, theoretical analysis based on inter-class correlation (ICC) suggests that our intuition in Fig. 1 holds statistically in CDFSL settings. Last but not the least, we claim that our proposed method can be easily added to existing transfer learning based methods to boost their performance.

Our major contributions are summarized as follows:

- To the best of our knowledge, our work is the first work that proposes to use supervised autoencoder framework to boost the model generalization capability under few-shot learning settings.
- We propose to take reconstructed images from autoencoder as noisy inputs and let the model further predict their labels, which proves to further enhance the model generalization capability. The two-step fine-tuning procedure that does reconstruction in novel classes better adapts model to the target domain.
- Extensive experiments across multiple benchmark datasets, various backbone architectures, and different loss function combinations demonstrate the efficacy and robustness of our proposed framework under cross-domain few-shot learning setting.

2. Related work

Few-Shot learning FSL aims at recognizing examples from novel categories with a limited number of labeled samples in each class. Meta-learning scheme for FSL receives much attention for its efficiency and simplicity. Existing meta-learning based methods can be classified into two general classes: the metric-based approaches [41, 14, 35, 36] that classify query images based on the similarity of feature embedding between query images and a few labeled images (support images), and the optimization-based approaches [11, 33, 43] that integrate the task-specific fine-tuning and pre-training into a single optimization framework. However, it is shown in [6, 16] that these SOTA methods for FSL underperform simple fine-tuning when the novel classes are from a different domain. Past works [8, 15] explore involving self-supervised learning scheme to obtain more diverse and transferable visual representations in few-shot learning. They fail to consider the domain-shift issue within the CDFSL settings.

Domain adaption The technique of domain adaption [44] is usually applied to solve the domain shift issue. It aims at learning a mapping from the source domain to the target domain so that the model trained on the source domain can be

applied to the target domain. However, there are some limitations of domain adaption that hinder its use in CDFSL. First, most domain adaption framework [12, 13, 18, 38] aims at learning the mapping under the same class. For example, learn the mapping from cartoon dogs to picture of an actual dogs. This does not fit into the FSL setting where the source and target domain have different classes. There are some existing works such as [10, 25] that consider the domain adaption technique under FSL settings. However, these approaches require a large set of unlabeled images in the target domain, which may be very difficult or even unrealistic in practice, e.g. X-ray and fMRI images.

Domain generalization Domain generalization methods differ from domain adaption in that they aim to generalize from a set of source domains to the target domains without accessing instances from the target domain during the training stage [39]. Past work to improve model generalization capability includes extracting domain-invariant features from various seen domains [27, 3, 24], decomposing the classifiers into domain-specific and domain-invariant components [19, 23], and augmenting the input data with adversarial learning [34, 42]. However, these methods require access to multiple source domains during training. Meta-learning methods achieve domain generalization by simulating testing scenarios in source domains during training, but they perform poorly when there is a domain-shift from source domain to target domain [6, 16].

Transfer learning Transfer learning is a more general term for methods in which different tasks or domains are involved. One traditional transfer learning approach is the simple fine-tuning. In the simple fine-tuning, a model is trained on the source dataset and the pre-trained model is then used as initialization to train the model on the target dataset. It is shown in [6, 16] that the simple fine-tuning can outperform all SOTA FSL methods under CDFSL setting. However, when the model overfits to the source domain, the fine-tuning performs worse than directly train the same model from random initialization. This is called negative transfer [7]. To avoid negative transfer and further improve the performance of simple fine-tuning under CDFSL, [26] proposes a batch spectral regularization (BSR) mechanism by penalizing the eigenvalues of the feature matrix.

3. Methodology

3.1. Preliminaries

Problem formulation In the cross-domain few-shot learning (CDFSL), we have a source domain \mathcal{T}_s and a target domain \mathcal{T}_t that have disjoint label sets. There exists a domain-shift between \mathcal{T}_s and \mathcal{T}_t [49]. The source domain has a large-scale labeled dataset \mathcal{D}_s while the target domain only has limited labeled images. Our method first pre-trains the model on the source dataset and then fine-tunes on the target

dataset. Each “N-way K-shot” classification task in target domain contains a support dataset \mathcal{D}_t^s and a query dataset \mathcal{D}_t^q . The support set contains N classes with K labeled images in each class and the query set contains images from the same N classes with Q unlabeled images in each class. The goal of CDFSL is to achieve a high classification accuracy on the query set \mathcal{D}_t^q when K is small.

Supervised autoencoder The autoencoder is a model that is usually used to obtain low-dimensional representations in an unsupervised manner. An autoencoder is composed of an encoder f_ϕ that encodes the input \mathbf{x} to its lower-dimensional representation $\tilde{\mathbf{x}} = f_\phi(\mathbf{x})$. Then, a decoder g_ψ decodes the representation $\tilde{\mathbf{x}}$ to $\hat{\mathbf{x}} = g_\psi(\tilde{\mathbf{x}})$ which is a reconstruction of input \mathbf{x} . The goal of the autoencoder is to minimize the difference between the input \mathbf{x} and its reconstruction $\hat{\mathbf{x}}$ and the reconstruction loss is formulated as

$$\mathcal{L}_{\text{REC}}(\phi, \psi; \mathbf{x}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_2. \quad (1)$$

When the labels of the inputs are available, the supervised autoencoder (SAE) [22] that jointly predicts the class label and reconstructs the input is proved to generalize well for downstream tasks. In the SAE, the representation $\tilde{\mathbf{x}}$ is fed into a classification module for label prediction and the loss function is

$$\mathcal{L}_{\text{SAE}}^{\lambda, \text{cls}}(\phi, \psi; \mathbf{x}, y) = \mathcal{L}_{\text{cls}}(\tilde{\mathbf{x}}, y) + \lambda \mathcal{L}_{\text{REC}}(\phi, \psi; \mathbf{x}) \quad (2)$$

where \mathcal{L}_{cls} is a loss function for classification and λ is a hyper-parameter that controls the reconstruction weight.

3.2. Overview

Under CDFSL setting, [16] shows that the traditional transfer learning based methods outperform all FSL methods. In the traditional transfer learning based method, a feature extractor is first pre-trained on the \mathcal{D}_s with sufficient labeled images by minimizing the classification loss $\mathcal{L}_{\text{cls-P}}$. Then the pre-trained feature extractor is fine-tuned on the target domain support set \mathcal{D}_t^s by minimizing the classification loss $\mathcal{L}_{\text{cls-F}}$. Note that the loss functions $\mathcal{L}_{\text{cls-P}}$ during the pre-training and $\mathcal{L}_{\text{cls-F}}$ during the fine-tuning may be different. Considering the superior performance of traditional transfer learning method on CDFSL, we use the transfer learning pipeline in our work. Motivated by the generalization capability of SAE and the generalization enhancement by feeding noisy inputs, we propose to boost the generalization capability of model via a noise-enhanced SAE (NSAE). To achieve this, we train a SAE that learns the feature space in generative and discriminative manners. NSAE not only predicts the class labels of the inputs but also predicts the labels of the “noisy” reconstructions. We also leverage the NSAE to perform domain adaption during the fine-tuning. Specifically, it is tuned to reconstruct target domain images before tuned to do classification task. An overview of our

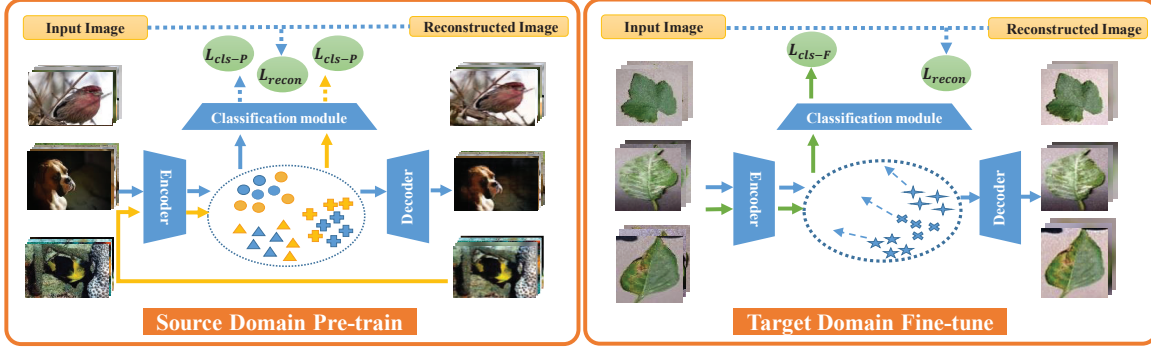


Figure 2: **An overview of the proposed pipeline.** A noise-enhanced supervised autoencoder (NSAE) is pre-trained with source dataset on the source domain to improve the generalization capability. The fine-tuning on the target domain is a two-step procedure that first performs reconstruction task on novel dataset, and then the encoder is fine-tuned for classification.

proposed pipeline is depicted in Fig. 2. A detailed explanation is given in the following sections.

3.3. Pre-train on the source domain

To borrow information from the source domain, the first step is to train a feature encoder on the large-scale source domain. Instead of training a single feature encoder, we propose to train a NSAE on the source domain. Let $\mathcal{D}_s = \{(\mathbf{x}_m^s, \mathbf{y}_m^s), m = 1, 2, \dots, M\}$ be the source dataset where M denotes the number of classes, and let f_ϕ and g_ψ be the encoder and decoder respectively. The input images are fed into f_ϕ to extract the feature representations which are fed into g_ψ to reconstruct the original inputs. Meanwhile, the feature representations are also fed into a classification module to predict the class labels of inputs. In our formulation, the reconstructed images are seen as “noisy” inputs which are further fed back into the encoder for classification. The NASE is trained to reconstruct the input images and predict the class labels of both original and reconstructed images. The loss function of NSAE during the pre-training is

$$\begin{aligned} \mathcal{L}_{\text{NSAE}}(\phi, \psi; \mathcal{D}_s) = & \frac{1}{M} \sum_{m=1}^M \mathcal{L}_{\text{SAE}}^{\lambda_1, \text{cls-P}}(\phi, \psi; \mathbf{x}_m^s, \mathbf{y}_m^s) \\ & + \frac{\lambda_2}{M} \sum_{m=1}^M \mathcal{L}_{\text{cls-P}}(\theta; f_\phi(\hat{\mathbf{x}}_m^s), \mathbf{y}_m^s) \end{aligned} \quad (3)$$

where $\mathcal{L}_{\text{cls-P}}$ is some classification loss and \mathcal{L}_{SAE} is given in (2). The second term is the classification loss of reconstructed images. The classification loss functions for the original inputs and the reconstructed images are the same. λ_1 and λ_2 are two hyper-parameters that control the weights of losses.

We show in the ablation study that the use of the classification loss based on the reconstructed images is indispensable which further improves the generalization capability of

the feature encoder.

3.4. Fine-tune on the target domain

The second stage is to fine-tune the pre-trained model on the target domain where only a very limited number of labeled examples are available. Based on the nature of our autoencoder architecture, we propose a two-step procedure for domain adaptation to the target domain.

Let $\mathcal{D}_t^s = \{(\mathbf{x}_{ij}, \mathbf{y}_{ij}); i = 1, 2, \dots, N, j = 1, \dots, K\}$ be the support set on the target domain. In the first step, we leverage the autoencoder architecture and propose to perform domain adaptation by reconstructing the support images for certain epochs. The model aims at minimizing reconstruction loss $\sum_{i,j} \mathcal{L}_{\text{REC}}(\phi, \psi; \mathbf{x}_{ij})$. In the second step of the fine-tuning, only the encoder is used to fine-tune on \mathcal{D}_t^s with the classification loss $\mathcal{L}_{\text{cls-F}}$. We show in the ablation study that such a two-step procedure works better than purely fine-tuning the encoder with $\mathcal{L}_{\text{cls-F}}$ on the target support set. We refer to them as one-step or two-step fine-tuning respectively in the following.

In traditional fine-tuning, all the parameters of the encoder or the first several layers of the encoder are fixed when the parameters of the classification module are updated. However, [16] shows that, under the CDFSL setting, the fine-tuned model can achieve better performance when the model is completely flexible. Therefore, we update all parameters of the model during the fine-tuning stage.

3.5. Choices of loss functions

The loss functions $\mathcal{L}_{\text{cls-P}}$ and $\mathcal{L}_{\text{cls-F}}$ are not specified in the description above. In fact, they can be any sensible loss functions for classification. In this paper, we study two loss functions for $\mathcal{L}_{\text{cls-P}}$ in pre-training stage, the first one is the cross entropy (CE) loss

$$\mathcal{L}_{\text{CE}}(\mathbf{W}; \mathbf{x}, \mathbf{y}) = -\log \left\{ \frac{\exp((\mathbf{W}\mathbf{x})_y)}{\sum_c \exp((\mathbf{W}\mathbf{x})_c)} \right\} \quad (4)$$

where \mathbf{W} is the parameters of the linear classifier and $(\cdot)_c$ means the c th element of the corresponding vector. The second one is the CE loss with batch spectral regularization (BSR) [26] that regularizes the singular values of the feature matrix in a batch. This classification loss is referred to as BSR loss and is given by

$$\mathcal{L}_{\text{BSR}}(\mathbf{W}) = \mathcal{L}_{\text{CE}}(\mathbf{W}) + \lambda \sum_{i=1}^b \sigma_i^2 \quad (5)$$

where $\sigma_1, \sigma_2, \dots, \sigma_b$ are singular values of the batch feature matrix.

In the second step of the fine-tuning stage, we consider the traditional fine-tuning and the distance-based fine-tuning. In traditional fine-tuning method, a linear classifier on top of the feature extractor is fine-tuned to minimize the CE loss. In the distance-based fine-tuning method, we follow the simple but effective distance-based classification method [35] in FSL, where the images are classified based on their similarities to the support images. To use distance-based loss function during the fine-tuning, at each iteration of the optimization, within each class of \mathcal{D}_t^s , we randomly split half of the images into a pseudo-support set $\mathcal{D}_t^{ps} = \{(\mathbf{x}_{ij}^s, y_{ij}^s), i = 1, 2, \dots, N, j = 1, 2, \dots, K/2\}$ and the rest to a pseudo-query set $\mathcal{D}_t^{pq} = \{(\mathbf{x}_{ij}^q, y_{ij}^q), i = 1, 2, \dots, N, j = 1, 2, \dots, K/2\}$. The feature embeddings of the pseudo-support set and the pseudo-query set based on the feature extractor f_ϕ is first obtained. Then the mean feature embeddings of the pseudo-support images in the same class

$$\mathbf{c}_i = \frac{K}{2} \sum_{j=1}^{K/2} f_\phi(\mathbf{x}_{ij}), \quad i = 1, 2, \dots, N \quad (6)$$

is used to represent the class and is called the class prototype. Given a distance function $d(\cdot, \cdot)$ and a pseudo-query image \mathbf{x} , the classification module produces a distribution over classes. The probability that \mathbf{x} belongs to class k is given as:

$$\mathbb{P}(y = k | \mathbf{x}) = \frac{\exp(-d(f_\phi(\mathbf{x}), \mathbf{c}_k))}{\sum_{k'} \exp(-d(f_\phi(\mathbf{x}), \mathbf{c}_{k'}))} \quad (7)$$

Since the true class labels of the pseudo-query images are known, the parameter ϕ can therefore be fine-tuned by maximizing the log-likelihood of the images in the query set, that is

$$\mathcal{L}_D(\phi) = \sum_{i,j} \log \mathbb{P}(y = y_{ij}^q | \mathbf{x}_{ij}^q). \quad (8)$$

It is shown in [35] that the distance-based classifier is effective. After the feature encoder is fine-tuned with the classification loss, we use the full support set to build the class prototypes and then classifies the query image into the class

that has the highest probability in (7). This is equivalent to classify the query image with the nearest neighbor classifier, the query image is classified to class k if it is closest to k th class prototype. We use cosine distance for $d(\cdot, \cdot)$ in our experiment.

The combination of the two loss functions for $\mathcal{L}_{\text{cls-P}}$ and the two loss functions for $\mathcal{L}_{\text{cls-F}}$ leads to 4 different loss functions respectively named as CE+CE, BSR+CE, CE+D, and BSR+D. The first acronym is referring to the loss function for $\mathcal{L}_{\text{cls-P}}$ and the second acronym is referring to the loss function for $\mathcal{L}_{\text{cls-F}}$.

4. Experiments

In this section, we demonstrate the efficacy of our proposed method for CDFSL on benchmark datasets via extensive experiments and ablation studies.

4.1. Experiment setting

Dataset Following the benchmark [16], we use *miniImageNet* as the source dataset, which is a subset of ILSVRC-2012 [32]. It contains 100 classes with 600 images in each class. Following the convention, the first 64 classes are used as the source domain images to pre-train the model in our experiment. To evaluate the generalization capability of our method, we use 8 different datasets as the target domains. The first four datasets are the benchmark datasets proposed in [16]. We refer to these four dataset as *CropDisease*, *EuroSAT*, *ISIC*, *ChestX* in the following, and the similarity of these datasets to mini-ImageNet decreases from left to right. We also include another four natural image datasets, *Car* [20], *CUB* [4, 45], *Plantae* [40], and *Places* [48] that are commonly used in CDFSL [39].

Evaluation protocol To make a fair comparison with existing methods for CDFSL, we evaluate the performance of the classifiers by simulating 600 independent 5-way few-shot classification tasks on each target domain dataset. For each task, we randomly sample 5 classes and within each class, we randomly select K images as the support set and 15 images as the query set. Following the benchmark [16], we let $K = 5, 20, 50$. In 50-shot classification, the *Car* dataset has only a few classes that have more than 50 images, so we do not consider this dataset; the *CUB* dataset has 144 out of 200 classes that have more than 60 images per class, so we sample from these 144 classes and use 10 images per class as query set for 5-way 50-shot evaluation. Then for each task, we fine-tune the pre-trained model on the support set and evaluate its performance on the query set. Transductive inference [16, 28] is used that the statistics of the query images are used in batch normalization. In total, the pre-trained model is fine-tuned and evaluated for 600 times under each experiment setting, and the average classification accuracy as well as 95% confidence interval on the query set is reported.

Network architecture To illustrate the effectiveness of the supervised autoencoder, we consider two commonly used encoder architectures in the experiment, namely Conv4 [41] and ResNet10 [16]. Besides the difference in network architecture, these two networks have different input sizes. We resize the source and target domain images to 84×84 for Conv4 and 224×224 for ResNet10. We design different decoder architectures for these two encoders. The decoders we designed consist of deconvolutional blocks, with each block containing 2D transposed convolution operator and ReLU activation, which expand the dimension of the feature map. To mirror the dimension of the output in the encoder, we set the hyperparameters in the 2D transposed convolution layer to be $kernelsize = 2$ and $stride = 2$. The architecture and layer specifications of the autoencoder can be found in Section 2 of the supplementary material.

Hyper-parameter settings All of our experiments are conducted in *pytorch* [29]. We use the same set of optimizers and hyper-parameters for all experiments regardless of model architecture and the target domain. Specially, in the pre-training stage, the model is trained from scratch, with a batch size of 64 for 400 epochs. We use combinations of random crop, random flip, and color jitter to augment the source dataset. We let $\lambda_1 = \lambda_2 = 1$ in (3) and let $\lambda = 0.001$ in (5). We optimize our model with stochastic gradient descent (SGD), with a learning rate of 10^{-3} , the momentum of 0.9, and weight decay of 5×10^{-4} . In the fine-tuning stage, we also use SGD optimization. In the first step, we use a learning rate of 10^{-3} and do reconstruction task for 30 epochs. In the second step, we use a learning rate of 10^{-2} , the momentum of 0.9, and weight decay of 10^{-3} and fine-tune for 200 epochs. In the distance-based fine-tuning, as pointed in Section 3.5, half of the support set is used as pseudo-support set and the other half is used as pseudo-query set. In the traditional fine-tuning, the batch size of 4 is used for 5 and 20 shot, and 16 for 50 shot.

Data augmentation & label propagation A simple but effective way in FSL is to supplement the small support set with hand-crafted data augmentation [26]. The operations such as random crop, random flip, and color jitter can be used to augment the dataset. We use the same combination of operations as shown in Table 1 in [26] when the data augmentation technique is used during the fine-tuning. For the distance-based fine-tuning, the order of data split and data augmentation leads to a difference in the training set. In our experiment, we first augment the support set and then randomly split the augmented images within the same class into a pseudo-support set and a pseudo-query set. For the traditional fine-tuning, we augment the support set and at each iteration of the fine-tuning, a random batch is selected to compute the gradient. To further improve the classification accuracy, a post-processing method called label propagation [26] is also used in our method. The label propa-

gation refines the predicted labels based on the similarities within the unlabeled query set.

4.2. Ablation study

To study the effectiveness of our proposed method, we conduct an ablation study under 5-way 5-shot setting on all 8 datasets with different architectures to show that

- (1) with the same classification loss combination, our proposed method boosts generalization capability and obtains consistently better performance on the target domain than the traditional transfer learning based methods for CDFSL;
- (2) the noise-enhancement that predicts class labels of reconstructed images is necessary and can greatly improve the generalization capability during pre-training;
- (3) the proposed two-step fine-tuning procedure achieves better domain adaption and leads to higher classification accuracy than the traditional one-step fine-tuning.

The ablation study is conducted using four kinds of combinations of the classification loss functions for pre-training and fine-tuning, i.e. CE+CE, BSR+CE, CE+D, and BSR+D. The average classification accuracy across 8 datasets is visualized in Fig. 3. The results based on two different encoder architectures, i.e. Conv4 and ResNet10, are respectively give in Fig. 3 (a) and Fig. 3 (b). The details of results of 8 datasets under different settings can be found in Section 3 in supplementary file.

To show (1), we compute the 5-way 5-shot classification accuracy when we only train a single encoder on the source domain and when we train a NSAE on the source

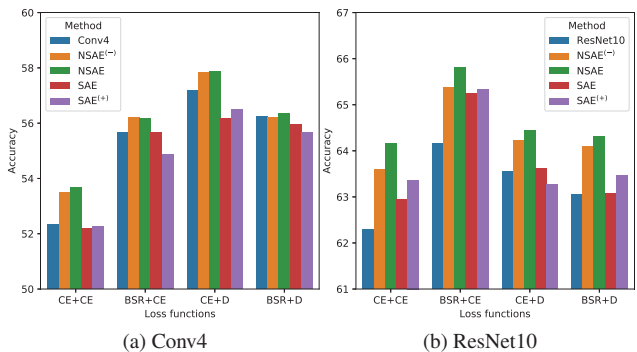


Figure 3: **Ablation study visualization.** The average 5-way 5-shot classification accuracy over 8 datasets when encoder is (a) Conv4 and (b) ResNet10. Within each plot, the bars are grouped by the classification loss functions during pre-training and fine-tuning on x-axis. Our proposed method NSAE is represented by the green bar.

domain. These two cases are labeled as ResNet10 (Conv4) and NSAE in Fig. 3. As is shown in the plot, our proposed method always has higher classification accuracy regardless of the encoder architecture and the classification functions.

To show (2), we compare our proposed method with two extreme cases. The first case is the SAE where we do not further feed in the reconstructed images for classification during the pre-training. The second case is the SAE^(*) where we double the weight on the classification loss of original images as if the auto-encoder works perfectly that the reconstructed images are identical to original images. As shown in the figure, the NSAE surpasses the other two variants under different settings. It suggests that the classification loss based on the reconstructed images is necessary. Without this loss, the two extreme cases that we compare with could even be worse than the traditional transfer learning based methods. We also compare our proposed method with that when hand-crafted noisy images are used, the results can be found in the supplementary material Section 4.

To show (3), we train two NSAEs with the same pre-training method and fine-tune the pre-trained autoencoder either with a one-step procedure or a two-step procedure as described in Section 3.4. These two cases are respectively named as NSAE(-) and NSAE. It is shown in the figure that the two-step fine-tune procedure also outperforms the one-step fine-tune procedure.

From the ablation study, we can also see that when the loss functions during the pre-training and fine-tuning are the same, the more complex encoder ResNet10 gives a higher classification accuracy compared with Conv4. When the pre-training classification loss is CE, using the distance-based loss function during the fine-tuning gives higher classification accuracy than using CE loss. However, when using classification loss BSR during pre-training, we get an opposite conclusion. Overall, using BSR as classification loss during pre-training and CE as classification loss during fine-tuning achieves the highest accuracy.

4.3. Generalization capability analysis

T-SNE visualization To qualitatively show the generalization capability of the feature encoder in our proposed method, we use t-SNE to visualize the feature embeddings of images from the source and target domain respectively in first row and second row in Fig. 4. In each plot, we randomly select 5 classes on each domain and visualize the features of all images in these classes based on different pre-trained encoders without fine-tuning. We use ResNet10 as encoder structure with CE (1st column) or BSR (3rd column) classification loss during the pre-training. Our proposed methods correspond to the figures in the even columns. As shown in the first row in Fig. 4, since there are enough training examples on the source domain, all models exhibit discriminative structures. The feature embeddings

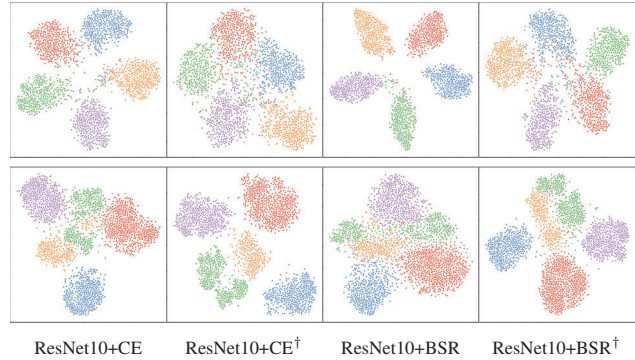


Figure 4: **Feature embedding visualization.** The t-SNE visualizations of the feature embeddings based on the pre-trained model on the source domain (1st row) and on the target CropDisease domain (2nd row). The method with † is our proposed feature extractor.

based on BSR loss are more centered than the CE loss, as the eigenvalues of the feature maps are regularized during the training. Moreover, the feature embeddings based on the NSAE losses have larger within-class variations and smaller class margins, as the model takes classification and reconstruction tasks at the same time. On the target domain, as shown in the second row in Fig. 4, we observe the opposite. In 1st and 3rd columns, features of different classes become confused with traditional pre-training. When the NSAE loss is used, the classes on the target domain becomes more separable. The within-class variations are smaller and the inter-class distance becomes larger. This suggests the better generalization capability of our proposed method.

Statistical analysis of discriminability Moreover, we also quantitatively measures the discriminability of the feature embeddings by the intra-class correlation (ICC). The ICC is defined as the ratio of inter-class variation and the intra-class variation. Therefore, the larger the ICC, the features in different classes are more separated or the features within the same classes are more concentrated. The details of the definition of ICC are in Section 1 in supplementary file. We compare the ICC of the features extracted from the traditionally pre-trained encoder and that based on our proposed NSAE without fine-tuning. We use two kinds of encoder, i.e. Conv4 and ResNet10, and two kinds of classification loss, i.e. CE and BSR, during pre-training. This leads to four combinations denoted as CE (Conv4), BSR (Conv4), CE (ResNet10), and BSR (ResNet10). We take the ratio of the ICCs of the traditionally trained method and that based on our proposed method. The results are given in Fig. 5(a). As shown in the figure, the ICC ratios are greater than 1 on the source domain (blue crosses) and smaller than 1 on the target domain (yellow stars) for all 4 scenarios. This suggests on the source domain, the feature extractor from our proposed method is not as discriminative as that trained

Table 1: **Comparison with SOTA methods.** The 5-way K-shot classification accuracy on 8 datasets with ResNet10 as the backbone. Our proposed method with CE+CE and BSR+CE losses are respectively denoted as “NASE[†]” and “NASE[‡]”. The (+) denotes that the data augmentation and label propagation techniques are used.

Methods	ISIC			EuroSAT			CropDisease			ChestX		
	5-shot	20-shot	50-shot	5-shot	20-shot	50-shot	5-shot	20-shot	50-shot	5-shot	20-shot	50-shot
Fine-tune[16]	48.11±0.64	59.31±0.48	66.48±0.56	79.08±0.61	87.64±0.47	90.89±0.36	89.25±0.51	95.51±0.31	97.68±0.21	25.97±0.41	31.32±0.45	35.49±0.45
NSAE [†]	54.05±0.63	66.17±0.59	71.32±0.61	83.96±0.57	92.38±0.33	95.42±0.34	93.14±0.47	98.30±0.19	99.25±0.14	27.10±0.44	35.20±0.48	38.95±0.70
BSR[26]	54.42±0.66	66.61±0.61	71.10±0.60	80.89±0.61	90.44±0.40	93.88±0.31	92.17±0.45	97.90±0.22	99.05±0.14	26.84±0.44	35.63±0.54	40.18±0.56
NSAE [‡]	55.27±0.62	67.28±0.61	72.90±0.55	84.33±0.55	92.34±0.35	95.00±0.26	93.31±0.42	98.33±0.18	99.29±0.14	27.30±0.42	35.70±0.47	38.52±0.71
LMM PQS[46]	51.88±0.60	64.88±0.58	69.46±0.58	86.30±0.53	92.59±0.31	94.16±0.28	93.52±0.39	97.60±0.23	98.24±0.17	26.10±0.44	32.58±0.47	38.22±0.52
NSAE [†] (+)	54.86±0.67	66.53±0.60	72.00±0.60	87.04±0.51	93.89±0.30	96.55±0.29	95.65±0.35	99.10±0.16	99.67±0.12	27.58±0.47	37.12±0.52	40.74±0.73
BSR(+)	56.82±0.68	67.31±0.57	72.33±0.58	85.97±0.52	93.73±0.29	96.07±0.30	95.97±0.33	99.10±0.12	99.66±0.07	28.50±0.48	36.95±0.52	42.32±0.53
NSAE [‡] (+)	56.85±0.67	67.45±0.60	73.00±0.56	87.53±0.50	94.21±0.29	96.50±0.29	96.09±0.35	99.20±0.14	99.70±0.09	28.73±0.45	36.14±0.50	41.80±0.72

Methods	Car			CUB			Plantae			Places		
	5-shot	20-shot	50-shot	5-shot	20-shot	50-shot	5-shot	20-shot	50-shot	5-shot	20-shot	50-shot
Fine-tune	52.08±0.74	79.27±0.63	–	64.14±0.77	84.43±0.65	89.61±0.55	59.27±0.70	75.35±0.68	81.76±0.56	70.06±0.74	80.96±0.65	84.79±0.58
NSAE [†]	54.91±0.74	79.68±0.54	–	68.51±0.76	85.22±0.56	89.42±0.62	59.55±0.74	75.70±0.64	82.42±0.55	71.02±0.72	82.70±0.58	85.90±0.59
BSR	57.49±0.72	81.56±0.78	–	69.38±0.76	85.84±0.79	90.91±0.56	61.07±0.76	77.20±0.90	82.16±0.59	71.09±0.68	81.76±0.81	85.67±0.57
NSAE [‡]	58.30±0.75	82.32±0.50	–	71.92±0.77	88.09±0.48	91.00±0.79	62.15±0.77	77.40±0.65	83.63±0.60	73.17±0.72	82.50±0.59	85.92±0.56
GNN-FT[39]	44.90±0.64	–	–	66.98±0.68	–	–	53.85±0.62	–	–	73.94±0.67	–	–
NSAE [†] (+)	55.51±0.73	83.17±0.56	–	69.96±0.80	89.01±0.54	93.11±0.64	61.71±0.79	78.58±0.64	84.64±0.76	71.86±0.72	83.24±0.58	86.22±0.70
BSR(+)	59.82±0.76	82.39±0.51	–	73.83±0.74	90.88±0.42	92.91±0.60	64.20±0.77	79.66±0.65	83.44±0.79	71.61±0.71	82.12±0.80	85.82±0.75
NSAE [‡] (+)	61.11±0.79	85.04±0.52	–	76.00±0.71	91.08±0.42	95.41±0.50	65.66±0.78	81.54±0.60	85.99±0.72	73.40±0.71	83.00±0.59	86.53±0.77

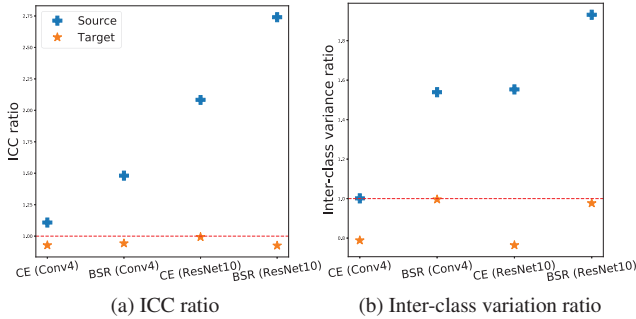


Figure 5: **ICC visualization.** The comparison of the ICC and the inter-class variance on the source domain and target datasets for different feature extractors.

with traditional methods. However, these feature extractors generalize better on the target domain. We similarly show the inter-class variations in Fig. 5(b). Our proposed method shows a larger inter-class variation on the target domain, suggesting that the classes are more separable.

4.4. Main results

Based on the ablation study, we use the combinations of **CE+CE** and **BSR+CE** as classification losses. We use [†] to denote method with **CE+CE** losses and [‡] to denote method with **BSR+CE** losses, and ResNet10 is used as feature encoder to compare with the SOTAs. Using traditional transfer learning, the CE+CE reduces to the “Fine-tune” method in [16] and BSR+CE reduces to the “BSR” method in [26]. Note that for these methods, since they only implement on ISIC, EuroSAT, CropDisease, and ChestX, we self-implement their models on Car, CUB, Plantae, and Places

datasets with their public codebase and report the results. To further improve the performance of our model, we also used data augmentation and label propagation techniques, denoted with (+) in Table 1. In Table 1, by comparing “Fine-tune” with “NASE[†]”, “BSR” with “NASE[‡]”, we can observe that our proposed method can improve on baselines by a large margin across different unseen domains for all shots. We attribute it to the great generalization ability of the pre-training mechanism and two-step domain adaption in the fine-tuning stage. Adding augmentation techniques can further improve the results and our proposed method performs favorably against other SOTAs across different unseen domains and evaluation settings.

5. Conclusion

In this work, we propose a novel method for improving the generalization capability of the transfer learning based methods for cross-domain few-shot learning(CDFSL). We propose to train a noise-enhanced supervised autoencoder instead of a simple feature extractor on the source domain. Theoretical analysis shows that NSAE can largely improve the generalization capability of the feature extractor. We also leverage the nature of the autoencoder and propose a two-step fine-tuning procedure that outperforms the past one-step fine-tune procedure. Extensive experiments and analysis demonstrate the efficacy and generalization of our method. Moreover, the formulation of NSAE makes it very easy to apply our proposed method to existing transfer learning based methods for CDFSL to further boost their performance.

References

- [1] Guozhong An. The effects of adding noise during backpropagation training on a generalization performance. *Neural computation*, 8(3):643–674, 1996. [2](#)
- [2] Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116, 1995. [2](#)
- [3] Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems*, 24:2178–2186, 2011. [3](#)
- [4] Steve Branson, Catherine Wah, Florian Schroff, Boris Babenko, Peter Welinder, Pietro Perona, and Serge Belongie. Visual recognition with humans in the loop. In *European Conference on Computer Vision*, pages 438–451. Springer, 2010. [5](#)
- [5] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019. [1](#)
- [6] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. *arXiv preprint arXiv:1904.04232*, 2019. [1](#), [2](#), [3](#)
- [7] Xinyang Chen, Sinan Wang, Bo Fu, Mingsheng Long, and Jianmin Wang. Catastrophic forgetting meets negative transfer: Batch spectral shrinkage for safe transfer learning. 2019. [3](#)
- [8] Zhixiang Chi, Yang Wang, Yuanhao Yu, and Jin Tang. Test-time fast adaptation for dynamic scene deblurring via meta-auxiliary learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9137–9146, 2021. [2](#)
- [9] 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. [1](#)
- [10] Nanqing Dong and Eric P Xing. Domain adaption in one-shot learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 573–588. Springer, 2018. [3](#)
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017. [1](#), [2](#)
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. [3](#)
- [13] 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. [3](#)
- [14] Victor Garcia and Joan Bruna. Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043*, 2017. [1](#), [2](#)
- [15] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8059–8068, 2019. [2](#)
- [16] Yunhui Guo, Noel C Codella, Leonid Karlinsky, James V Codella, John R Smith, Kate Saenko, Tajana Rosing, and Rogerio Feris. A broader study of cross-domain few-shot learning. *ECCV*, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#)
- [18] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. Progressive domain adaptation for object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 749–757, 2020. [3](#)
- [19] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the damage of dataset bias. In *European Conference on Computer Vision*, pages 158–171. Springer, 2012. [3](#)
- [20] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. [5](#)
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [1](#)
- [22] Lei Le, Andrew Patterson, and Martha White. Supervised autoencoders: Improving generalization performance with unsupervised regularizers. *Advances in neural information processing systems*, 31:107–117, 2018. [2](#), [3](#)
- [23] 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. 3
- [24] 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. 3
- [25] Xiang Li, Xiao-Dong Jia, Wei Zhang, Hui Ma, Zhong Luo, and Xu Li. Intelligent cross-machine fault diagnosis approach with deep auto-encoder and domain adaptation. *Neurocomputing*, 383:235–247, 2020. 3
- [26] Bingyu Liu, Zhen Zhao, Zhenpeng Li, Jianan Jiang, Yuhong Guo, Haifeng Shen, and Jieping Ye. Feature transformation ensemble model with batch spectral regularization for cross-domain few-shot classification. *arXiv preprint arXiv:2005.08463*, 2020. 2, 3, 5, 6, 8
- [27] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18. PMLR, 2013. 3
- [28] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 5
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019. 6
- [30] Niamul Quader, Md Mafijul Islam Bhuiyan, Juwei Lu, Peng Dai, and Wei Li. Weight excitation: Built-in attention mechanisms in convolutional neural networks. In *European Conference on Computer Vision*, pages 87–103. Springer, 2020. 1
- [31] Russell Reed and Robert J MarksII. *Neural smithing: supervised learning in feedforward artificial neural networks*. Mit Press, 1999. 2
- [32] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 5
- [33] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018. 1, 2
- [34] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018. 3
- [35] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017. 1, 2, 5
- [36] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208, 2018. 1, 2
- [37] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019. 1
- [38] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7472–7481, 2018. 3
- [39] Hung-Yu Tseng, Hsin-Ying Lee, Jia-Bin Huang, and Ming-Hsuan Yang. Cross-domain few-shot classification via learned feature-wise transformation. *arXiv preprint arXiv:2001.08735*, 2020. 2, 3, 5, 8
- [40] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018. 5
- [41] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016. 1, 2, 6
- [42] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *arXiv preprint arXiv:1805.12018*, 2018. 3
- [43] Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J Lim. Multimodal model-agnostic meta-learning via task-aware modulation. *arXiv preprint arXiv:1910.13616*, 2019. 1, 2
- [44] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018. 2
- [45] Peter Welinder and Pietro Perona. Online crowdsourcing: rating annotators and obtaining cost-effective la-

- bels. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*, pages 25–32. IEEE, 2010. 5
- [46] Jia-Fong Yeh, Hsin-Ying Lee, Bing-Chen Tsai, Yi-Rong Chen, Ping-Chia Huang, and Winston H Hsu. Large margin mechanism and pseudo query set on cross-domain few-shot learning. *arXiv preprint arXiv:2005.09218*, 2020. 8
- [47] An Zhao, Mingyu Ding, Zhiwu Lu, Tao Xiang, Yulei Niu, Jiechao Guan, and Ji-Rong Wen. Domain-adaptive few-shot learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1390–1399, 2021. 2
- [48] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017. 5
- [49] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020. 1, 3