

REVISITING FINE-TUNING FOR FEW-SHOT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Few-shot learning is the process of learning novel classes using only a few examples and it remains a challenging task in machine learning. Many sophisticated few-shot learning algorithms have been proposed based on the notion that networks can easily overfit to novel examples if they are simply fine-tuned using only a few examples. In this study, we show that in the commonly used low-resolution *mini*-ImageNet dataset, the fine-tuning method achieves higher accuracy than common few-shot learning algorithms in the 1-shot task and nearly the same accuracy as that of the state-of-the-art algorithm in the 5-shot task. We then evaluate our method with more practical tasks, namely the high-resolution single-domain and cross-domain tasks. With both tasks, we show that our method achieves higher accuracy than common few-shot learning algorithms. We further analyze the experimental results and show that: 1) the retraining process can be stabilized by employing a low learning rate, 2) using adaptive gradient optimizers during fine-tuning can increase test accuracy, and 3) test accuracy can be improved by updating the entire network when a large domain-shift exists between base and novel classes.

1 INTRODUCTION

Previous studies have shown that high image classification performance can be achieved by using deep networks and big datasets (Krizhevsky et al., 2012; Simonyan & Zisserman, 2015; He et al., 2016; Szegedy et al., 2015). However, the performances of these algorithms rely heavily on extensive manually annotated images, and considerable cost is often incurred in preparing these datasets. To avoid this problem, few-shot learning, which is a task of learning novel classes using only a few examples, has been actively researched.

However, few-shot learning remains a considerably challenging task in machine learning, and classification accuracy in few-shot tasks is much lower than that of the *many*-shot regime. This is because a network pretrained using base classes must adapt to novel classes using only a few examples. The simplest means of overcoming this difficulty is to fine-tune the network using novel classes. However, the number of trainable parameters of deep networks is so large that we believe that networks can easily overfit to novel classes if we simply fine-tune the networks using only a few examples. For example, the number of trainable parameters in the ResNet-152 (He et al., 2016) is approximately 60 M, which is much greater than the number of novel examples (e.g., 25 for 5-way 5-shot learning), and this leads us to the idea of overfitting. Using various sophisticated methods, numerous studies have been conducted to prevent networks from overfitting. However, the performance of a naive fine-tuning method has not been well investigated, and Chen et al. (2019) has pointed out that performance of this method had been underestimated in previous studies. Therefore, in this study, we analyze the performance of a fine-tuning method and show that it can achieve higher classification accuracy than common few-shot learning methods and, in some cases, can achieve an accuracy approximating that of the state-of-the-art algorithm. We also experimentally show that: 1) a low learning rate stabilizes the retraining process, 2) using an adaptive gradient optimizer when fine-tuning the network increases test accuracy, and 3) updating the entire network increases test accuracy when a large domain shift occurs between base and novel classes.

To evaluate accuracy in few-shot image classification tasks, the *mini*-ImageNet dataset (Vinyals et al., 2016) has been used in many previous studies. This is a subset of the ImageNet dataset (Deng et al., 2009) in which each image is resized to 84×84 to reduce computational cost. Recently, Chen et al. (2019) evaluated few-shot learning methods in more practical datasets, namely,

the high-resolution *mini*-ImageNet dataset and cross-domain dataset. Both datasets contain higher-resolution images than the original *mini*-ImageNet dataset, and the cross-domain dataset represents a greater challenge because base and novel classes are sampled from different datasets. Thus, a larger domain shift occurs between these classes. In this study, we evaluate the performance of our method using the high-resolution *mini*-ImageNet dataset (high-resolution single-domain task) and cross-domain dataset (cross-domain task) as well as the common low-resolution *mini*-ImageNet dataset (low-resolution single-domain task). Details of these datasets are provided in Section 2.3.

The main contributions of this study are as follows:

- 1) We show that in the common low-resolution single-domain task, our fine-tuning method achieves higher accuracy than common few-shot learning algorithms in the 1-shot task and nearly the same accuracy as that of the state-of-the-art method in the 5-shot task. We also show that our method achieves higher accuracy than common few-shot learning methods both in the high-resolution single-domain and cross-domain tasks. Note that we do not compare the performance of our method with the state-of-the-art algorithm in the high-resolution single-domain and cross-domain tasks because the performances for these tasks are not reported in the corresponding papers.
- 2) We further analyze the experimental results and show that a low learning rate stabilizes the re-learning process, that test accuracy can be increased by using an adaptive gradient optimizer such as the Adam optimizer, and that updating the entire network can increase test accuracy when a large domain shift occurs.

2 OVERVIEW OF FEW-SHOT LEARNING

2.1 NOTATION

Few-shot learning is a task of learning novel classes using only a few labeled examples. This task is also called N -way K -shot learning, where N denotes the number of novel classes and K is the number of labeled examples per class. We focus on the 5-way learning task such as in previous studies (Chen et al., 2019; Schwartz et al., 2018). Labeled and unlabeled examples of novel classes are called support and query sets, respectively. A network is pretrained using base classes, which contain numerous labeled examples. Base and novel classes are mutually exclusive. Base classes are used for pretraining, and novel classes are used for retraining and testing. Validation classes are used to determine a learning rate and the number of epochs required to retrain the network.

2.2 FEW-SHOT LEARNING ALGORITHMS

To date, numerous few-shot learning algorithms have been proposed, and these methods can be roughly classified into three categories: learning discriminative embedding using metric-based classification, learning to learn novel classes, and data-augmentation using synthetic data.

Metric-based Approach

Metric-learning approaches such as MatchingNet (Vinyals et al., 2016) and ProtoNet (Snell et al., 2017) tackle few-shot classification tasks by training an embedding function and applying a differentiable nearest-neighbor method to the feature space using the Euclidean metric. RelationNet (Sung et al., 2018) was developed to replace the nearest-neighbor method with a trainable metric using convolutional and fully connected (FC) layers and has achieved higher few-shot classification accuracy. Qi et al. (2018) proposed a method called weight imprinting. They showed that including normalized feature vectors of novel classes in the final layer weight provides effective initialization for novel classes. We use the weight-imprinting method to initialize the last FC layer before fine-tuning the network. These conventional methods successfully retrain networks using novel classes while preventing overfitting, but we show that few-shot classification performance can be further improved by fine-tuning networks.

Meta-learning-based Approach

Meta-learning-based approaches address the few-shot learning problem by training networks to *learn to learn* novel classes. Ravi & Larochelle (2017) focused on the similarity between gradient descent methods and long short-term memory (LSTM) (Hochreiter & Schmidhuber, 1997), and they achieved a fast adaptation to novel classes by using LSTM to update network weights. Finn et al.

(2017) proposed a method to train a network to obtain easily adaptable parameters so that the network can adapt to novel classes by means of a few gradient steps using only a few examples. In these algorithms, networks are explicitly trained to learn how to adapt to novel classes. However, we show that networks pretrained without explicit meta-learning methods can also learn novel classes and achieve high few-shot classification accuracy.

Data-augmentation-based Approach

Data-augmentation-based approaches overcome data deficiencies by generating synthetic examples of novel classes. Some methods synthesize examples of novel classes by applying within-class differences of base classes to real examples of novel classes (Hariharan & Girshick, 2017; Schwartz et al., 2018). Wang et al. (2018) integrated a feature generator using a few-shot learning process and succeeded in generating synthetic data using only a few novel examples. These methods succeeded in improving the performance of few-shot learning by using synthetic examples for retraining. Nevertheless, we show that networks can adapt to novel classes by using only naive data-augmentation methods such as image flipping and image jittering.

2.3 DATASETS FOR FEW-SHOT LEARNING

The *mini*-ImageNet dataset is a well-known dataset used to evaluate few-shot learning methods. The dataset was first proposed by Vinyals et al. (2016), but the train/validation/test split proposed by Ravi & Larochelle (2017) is often used instead. Therefore, we used this split in this study. This dataset is a subset of the ImageNet dataset (Deng et al., 2009) and contains 100 classes with 600 examples of each class. The classes are split into 64 base, 16 validation, and 20 novel classes. Images in this dataset are resized to 84×84 to reduce computational cost.

Recently, Chen et al. (2019) used a higher-resolution *mini*-ImageNet dataset with an image resolution of 224×224 to employ deeper networks. They also revealed that the domain shift that occurs between base and novel classes in the *mini*-ImageNet dataset is small because the classes are sampled in the same dataset. The authors proposed the cross-domain dataset, which has a larger domain shift between these classes. In this dataset, the whole *mini*-ImageNet dataset is used as a set of base classes, and randomly sampled 50 and 50 classes from the CUB-200-2011 dataset (Wah et al., 2011) are used as validation and novel classes, respectively. These datasets are more practical because they use high-resolution images, and the cross-domain dataset is more challenging because the domain shift that occurs between base and novel classes is larger. Therefore, we use the high-resolution *mini*-ImageNet dataset and cross-domain dataset for evaluation as well as the common low-resolution *mini*-ImageNet dataset.

3 EXPERIMENTS

3.1 NETWORK STRUCTURE

The ResNet-18/34/50/101/152 (He et al., 2016) and VGG-16 (Simonyan & Zisserman, 2015) without FC layers are used as feature extractors in this study. Note that the last MaxPool2d layer of the VGG-16 is replaced by the GlobalAveragePool2d layer to support different resolutions of input images. We also use the simple classifier (i.e., common FC layer) and the normalized classifier. The technique known as weight imprinting (Qi et al., 2018) is used in the normalized classifier; the normalized classifier is illustrated in Figure 1. Before fine-tuning the normalized network for novel classes, we initialize classifier weight W by deleting the weight and inserting columns for novel classes, as shown in Figure 1. Regarding the simple classifier, the initial weight for novel classes can be obtained by applying the multi-class linear SVM to feature vectors of novel classes.

3.2 DATASET

We evaluated our method using the low-resolution *mini*-ImageNet dataset as a common evaluation dataset. In addition, we used the high-resolution *mini*-ImageNet and cross-domain datasets as more practical datasets. Details of these datasets are provided in Section 2.3. In this study, we identify tasks that use these datasets as low-resolution single-domain task, the high-resolution single-domain task, and cross-domain tasks.




Figure 1: For the normalized classifier, we use the technique called weight imprinting proposed by Qi et al. (2018). Each column $w_i \in \mathbb{R}^d$ of the classifier weight is normalized so that w_i has a norm of 1 (i.e., $\|w_i\| = 1$), and classification is performed by taking the inner products between w_i and normalized feature vector $\hat{z} \in \mathbb{R}^d$. Note that variable d is the dimension of the feature space. Before the network is fine-tuned, the initial weight for a novel class can be obtained by including feature vector \hat{z} of the novel class in the classifier weight. When multiple novel examples per class are available (i.e., K -shot learning with $K > 1$), the initial weight for the novel class can be obtained by normalizing class mean $1/K \sum_{j=1}^K \hat{z}_j$ again. Because the output range of $W\hat{z}$ is $[-1, 1]$, ensuring that the probability of the correct label approximates 1 using softmax activation is difficult. This problem can be avoided by applying scale factor $s \in \mathbb{R}$ to the output, as discussed by Qi et al. (2018).

3.3 PRETRAINING THE NETWORK

The networks were pretrained by using the base classes of the datasets for 600 epochs. We used the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 0.001 in the same manner as Chen et al. (2019). These parameters for pretraining are normally optimized using validation classes, but we fixed these parameters to reduce computational cost. Input images were preprocessed by random-resized cropping with a size of 224×224 . We also performed color jittering and random-horizontal flipping, and we subtracted channel-wise means of the ImageNet dataset (0.485, 0.456, 0.406). In addition, division by channel-wise standard deviations of the ImageNet dataset (0.229, 0.224, 0.225) was performed in the same manner as Chen et al. (2019). Note that in the low-resolution single-domain task, the size of the random-resized cropping was set to 84×84 .

3.4 UPDATING THE NETWORK FOR NOVEL CLASSES

In this study, we compared three fine-tuning methods in which: 1) the entire network is updated, 2) the classifier weight and batch-normalization (BN) statistics are updated, and 3) only the classifier weight is updated. The third method is a common fine-tuning method to prevent overfitting. The second method is based on a previous study (Noguchi & Harada, 2019) that successfully fine-tuned an image generator to a novel class without overfitting by updating only the BN statistics (i.e., γ and β of BN layers). A similar approach is known as meta-transfer learning (MTL) (Sun et al., 2019). The authors who proposed MTL showed that updating only scales and biases of network parameters prevents them from overfitting while achieving efficient adaptation to unseen tasks. Although the methods proposed by Noguchi & Harada (2019) and Sun et al. (2019) presented similar ideas, we chose the former because of its simplicity in implementation.

Initial classifier weights for novel classes were obtained before we fine-tuned the networks, as discussed in Section 3.1. The networks were retrained with mini-batch-based learning with a batch size of NK in the N -way K -shot learning scenario. The learning rate and number of epochs for fine-tuning were determined by using validation classes. We evaluated few-shot classification accuracy by calculating the mean accuracy of 600 trials using randomly sampled classes and examples in the novel classes. We also calculated the 95% confidence interval of the mean accuracy. In the validation, test, and network initialization phases for the novel classes, input images were preprocessed by resizing to 256×256 , center-cropping to a size of 224×224 , subtracting channel-wise means of the ImageNet dataset (0.485, 0.456, 0.406), and dividing by channel-wise standard-deviations of the ImageNet dataset (0.229, 0.224, 0.225) in the same manner as Chen et al. (2019). The input preprocessing for fine-tuning phase was the same as discussed in Section 3.3.

Table 1: Performance of our method in the 5-way low-resolution single-domain task. “Normalized” and “Simple” mean that the normalized and simple classifiers are used, respectively. “All”, “BN & FC”, and “FC” mean the following: the entire network was updated; the BN and FC layer were updated; only the FC layers were updated. “w/o FT” refers to performance without fine-tuning using novel classes. Values with the † mark refer to classification accuracy without fine-tuning, as validation accuracy was not increased by fine-tuning the network. The * mark means that the classification accuracy for novel classes was not available because the loss value did not decrease in the pretraining phase. The - mark means that we did not conduct an experiment because the network did not have BN layers.

	Normalized All	Normalized BN & FC	Normalized FC	Normalized w/o FT	Simple All	Simple BN & FC	Simple FC	Simple w/o FT
1-shot ResNet-18	48.38 ± 0.68†	48.38 ± 0.68†	48.38 ± 0.68†	48.38 ± 0.68	49.27 ± 0.68†	49.27 ± 0.68†	49.27 ± 0.68†	49.27 ± 0.68
1-shot ResNet-34	48.41 ± 0.69†	48.41 ± 0.69†	48.41 ± 0.69†	48.41 ± 0.69	48.14 ± 0.67†	48.14 ± 0.67†	48.14 ± 0.67†	48.14 ± 0.67
1-shot ResNet-50	50.58 ± 0.74†	50.58 ± 0.74†	50.58 ± 0.74†	50.58 ± 0.74	49.48 ± 0.66†	49.48 ± 0.66†	49.48 ± 0.66†	49.48 ± 0.66
1-shot ResNet-101	48.83 ± 0.66†	48.83 ± 0.66†	48.83 ± 0.66†	48.83 ± 0.66	46.78 ± 0.67†	46.78 ± 0.67†	46.78 ± 0.67†	46.78 ± 0.67
1-shot ResNet-152	48.73 ± 0.70†	48.73 ± 0.70†	48.73 ± 0.70†	48.73 ± 0.70	49.61 ± 0.65†	49.61 ± 0.65†	49.61 ± 0.65†	49.61 ± 0.65
1-shot VGG-16	54.73 ± 0.71	-	54.90 ± 0.66	49.12 ± 0.71	*	-	*	*
5-shot ResNet-18	70.81 ± 0.54	70.08 ± 0.58	64.55 ± 0.58†	64.55 ± 0.58	69.50 ± 0.58	69.55 ± 0.54	68.32 ± 0.56	67.19 ± 0.56
5-shot ResNet-34	71.90 ± 0.55	70.61 ± 0.54	65.17 ± 0.57†	65.17 ± 0.57	69.48 ± 0.61	69.09 ± 0.56	67.92 ± 0.58	67.23 ± 0.60
5-shot ResNet-50	70.80 ± 0.57	69.55 ± 0.55	66.22 ± 0.57†	66.22 ± 0.57	68.92 ± 0.58†	68.92 ± 0.58†	68.92 ± 0.58†	68.92 ± 0.58
5-shot ResNet-101	69.65 ± 0.56	68.81 ± 0.56	65.02 ± 0.54†	65.02 ± 0.54	65.92 ± 0.60†	65.92 ± 0.60†	65.92 ± 0.60†	65.92 ± 0.60
5-shot ResNet-152	67.99 ± 0.53	66.62 ± 0.56	64.48 ± 0.54†	64.48 ± 0.54	70.24 ± 0.56†	70.24 ± 0.56†	70.24 ± 0.56†	70.24 ± 0.56
5-shot VGG-16	74.50 ± 0.50	-	69.92 ± 0.50	66.86 ± 0.55	*	-	*	*

Table 2: Performance of our method in the 5-way high-resolution single-domain task. “Normalized” and “Simple” mean that the normalized and simple classifiers were used, respectively. “All”, “BN & FC”, and “FC” mean the following: the entire network was updated; the BN and FC layer were updated; only the FC layers were updated. “w/o FT” refers to performance without fine-tuning using novel classes. Values with the † mark refer to classification accuracy without fine-tuning, as validation accuracy was not increased by fine-tuning the network. The * mark means that classification accuracy for novel classes was not available because the loss value did not decrease in the pretraining phase. The - mark means that we did not conduct an experiment because the network did not have BN layers.

	Normalized All	Normalized BN & FC	Normalized FC	Normalized w/o FT	Simple All	Simple BN & FC	Simple FC	Simple w/o FT
1-shot ResNet-18	56.28 ± 0.75†	56.28 ± 0.75†	56.28 ± 0.75†	56.28 ± 0.75	57.06 ± 0.67†	57.06 ± 0.67†	57.06 ± 0.67†	57.06 ± 0.67
1-shot ResNet-34	56.60 ± 0.73†	56.60 ± 0.73†	56.60 ± 0.73†	56.60 ± 0.73	56.83 ± 0.70†	56.83 ± 0.70†	56.83 ± 0.70†	56.83 ± 0.70
1-shot ResNet-50	58.05 ± 0.73†	58.05 ± 0.73†	58.05 ± 0.73†	58.05 ± 0.73	58.33 ± 0.70†	58.33 ± 0.70†	58.33 ± 0.70†	58.33 ± 0.70
1-shot ResNet-101	54.03 ± 0.71†	54.03 ± 0.71†	54.03 ± 0.71†	54.03 ± 0.71	51.96 ± 0.68†	51.96 ± 0.68†	51.96 ± 0.68†	51.96 ± 0.68
1-shot ResNet-152	56.73 ± 0.77†	56.73 ± 0.77†	56.73 ± 0.77†	56.73 ± 0.77	57.61 ± 0.71†	57.61 ± 0.71†	57.61 ± 0.71†	57.61 ± 0.71
1-shot VGG-16	59.40 ± 0.73	-	60.88 ± 0.71	57.15 ± 0.72	*	-	*	*
5-shot ResNet-18	78.03 ± 0.51	76.87 ± 0.52	72.94 ± 0.55†	72.94 ± 0.55	76.58 ± 0.54	77.20 ± 0.49	75.80 ± 0.50	75.08 ± 0.51
5-shot ResNet-34	77.73 ± 0.50	76.85 ± 0.50	71.99 ± 0.58†	71.99 ± 0.58	77.93 ± 0.53	76.15 ± 0.47	74.88 ± 0.52†	74.88 ± 0.52
5-shot ResNet-50	79.82 ± 0.49	79.20 ± 0.52	74.79 ± 0.54†	74.79 ± 0.54	77.94 ± 0.53†	77.94 ± 0.53†	77.94 ± 0.53†	77.94 ± 0.53
5-shot ResNet-101	78.08 ± 0.50	77.82 ± 0.48	73.33 ± 0.50†	71.18 ± 0.55	76.62 ± 0.52	77.41 ± 0.51	78.57 ± 0.48	72.32 ± 0.53
5-shot ResNet-152	78.98 ± 0.49	78.88 ± 0.46	74.57 ± 0.55†	74.57 ± 0.55	78.55 ± 0.48†	78.55 ± 0.48†	78.55 ± 0.48†	78.55 ± 0.48
5-shot VGG-16	78.80 ± 0.50	-	77.48 ± 0.49	75.78 ± 0.53	*	-	*	*

3.5 RESULTS

Few-shot classification accuracies for the low-resolution single-domain task, high-resolution single-domain task, and cross-domain task are listed in Tables 1, 2, and 3, respectively.

Table 1 shows that the classification accuracy could be increased by approximately 6% when the VGG-16 and normalized classifier were used in the 1-shot learning task. However, the accuracy could not be further improved by updating the entire network in the 1-shot learning task. However, classification accuracy could be further improved by updating the entire network in the 5-shot task. We assume that this was because the within-class difference could be reduced by fine-tuning the feature extractor when multiple novel examples were available.

By comparing the results for the high-resolution single-domain (Table 2) and low-resolution (Table 1) tasks, it could be argued that the robustness against low-resolution inputs differs depending on the feature extractor. For example, by comparing the results for 5-shot “Normalized all” in Table 1 and 2, we can see that the classification accuracy of the ResNet-152 decreased by 11.0% whereas that of the VGG-16 decreased by only 4.3%. This implies that the robustness against low-resolution inputs should also be considered and that evaluating only few-shot learning performance is difficult. Although the low-resolution *mini*-ImageNet dataset is extremely useful for doing fast experiments, we must reconsider the validity of the dataset for evaluation of few-shot learning performance.

Table 3: Performance of our method in the 5-way cross-domain task. “Normalized” and “Simple” mean that the normalized and simple classifiers were used, respectively. “All”, “BN & FC”, and “FC” mean the following: the entire network was updated; the BN and FC layer were updated; only the FC layers were updated. “w/o FT” refers to performance without fine-tuning using novel classes. Values with the † mark refer to classification accuracy without fine-tuning, as validation accuracy was not increased by fine-tuning the network. The * mark means that classification accuracy for novel classes was not available because the loss value did not decrease in the pretraining phase. The - mark means that we did not conduct an experiment because the network did not have BN layers.

	Normalized All	Normalized BN & FC	Normalized FC	Normalized w/o FT	Simple All	Simple BN & FC	Simple FC	Simple w/o FT
5-shot ResNet-18	71.83 ± 0.57	68.19 ± 0.65	60.26 ± 0.62	62.11 ± 0.68	68.67 ± 0.66	65.59 ± 0.62	65.51 ± 0.65	64.96 ± 0.68
5-shot ResNet-34	71.84 ± 0.61	66.93 ± 0.64	60.22 ± 0.67†	60.22 ± 0.67	71.43 ± 0.57	68.36 ± 0.59	66.62 ± 0.58	67.76 ± 0.62
5-shot ResNet-50	74.88 ± 0.58	70.89 ± 0.63	61.45 ± 0.68†	61.45 ± 0.68	69.27 ± 0.63	67.59 ± 0.62	67.77 ± 0.66	62.30 ± 0.63
5-shot ResNet-101	71.95 ± 0.62	68.58 ± 0.65	60.94 ± 0.57	53.78 ± 0.60	65.32 ± 0.59	63.93 ± 0.61	66.44 ± 0.62	49.86 ± 0.58
5-shot ResNet-152	73.56 ± 0.53	70.16 ± 0.62	61.59 ± 0.55	61.54 ± 0.66	66.90 ± 0.64	67.20 ± 0.64	68.30 ± 0.62	53.45 ± 0.63
5-shot VGG-16	71.92 ± 0.64	-	65.56 ± 0.57	59.84 ± 0.64	*	-	*	*

Table 4: Comparison between our method and conventional methods. We show several results from different networks in Section 3.5, and therefore in this table show the highest accuracy for each task. More specifically, we use the results from “VGG-16 Normalized FC”, “VGG-16 Normalized All”, “VGG-16 Normalized FC”, “ResNet-50 Normalized All”, and “ResNet-50 Normalized All” from left to right in this table. Values with the ‡ marks were reported by Chen et al. (2019), and other values were reported in the original studies. The - mark means that the classification accuracy for the task was not reported.

	Low-resolution Single-domain		High-resolution Single-domain		Cross-domain
	1-shot	5-shot	1-shot	5-shot	
Fine-tune (Ours)	54.90 ± 0.66	74.50 ± 0.50	60.88 ± 0.71	79.82 ± 0.49	74.88 ± 0.58
Baseline (Chen et al., 2019)	42.11 ± 0.71	62.53 ± 0.69	52.37 ± 0.79‡	74.69 ± 0.64‡	65.57 ± 0.70‡
Baseline++ (Chen et al., 2019)	48.24 ± 0.75	66.43 ± 0.63	53.97 ± 0.79‡	75.90 ± 0.61‡	62.04 ± 0.76‡
MatchingNet (Vinyals et al., 2016)	46.6	60.0	54.49 ± 0.81‡	68.88 ± 0.69‡	53.07 ± 0.74‡
ProtoNet (Snell et al., 2017)	49.42 ± 0.78	68.20 ± 0.66	54.16 ± 0.82‡	74.65 ± 0.64‡	62.02 ± 0.70‡
MAML (Finn et al., 2017)	48.70 ± 1.75	63.15 ± 0.91	54.69 ± 0.89‡	66.62 ± 0.83‡	51.34 ± 0.72‡
RelationNet (Sung et al., 2018)	50.44 ± 0.82	65.32 ± 0.70	53.48 ± 0.86‡	70.20 ± 0.66‡	57.71 ± 0.73‡
MTL (Sun et al., 2019)	61.2 ± 1.8	75.5 ± 0.8	-	-	-
Delta Encoder (Schwartz et al., 2018)	59.9	69.7	-	-	-

The comparison of the results from the cross-domain task (Table 3) and high-resolution single-domain task shows that the performance decreased in the cross-domain task. This can be explained by the larger domain shift that occurs between base and novel classes, as indicated by Chen et al. (2019). In addition, the difference in classification accuracy between the cross-domain task and high-resolution single-domain task was decreased by fine-tuning the entire network. For example, in the 5-shot learning task using the VGG-16, the difference in classification accuracy between the single-domain and cross-domain tasks was 15.9% without fine-tuning, but it could be reduced to 6.9% by fine-tuning the entire network. This means that the network could adapt to a large domain shift by having the entire network updated. We discuss this further in greater detail in Section 3.7.

3.6 COMPARATIVE EVALUATION

Table 4 shows comparative results of ours and previous methods. Note that we use the best result for each task as given in Section 3.5 because we obtained several results from different networks. In the 1-shot low-resolution single-domain task, the classification accuracy was lower than that of the state-of-the-art algorithm, but it was still higher than those of other common few-shot learning methods such as MatchingNet and ProtoNet. It is interesting to note that our method achieved nearly the same classification accuracy as that of the state-of-the-art method in the 5-shot task. The reason for the higher performance in the 5-shot task may be that the within-class variance could be reduced by fine-tuning the entire network using several examples per class.

In addition, we achieved higher classification accuracy than the reported values of conventional methods both in the high-resolution single-domain and cross-domain tasks. The difference in classification accuracy between the 5-shot high-resolution single-domain and cross-domain tasks was only approximately 5% with our method, whereas the performance of the conventional methods decreased by more than 10% in the cross-domain task. This shows that our method successfully




Figure 2: Classification accuracy for validation classes with different learning rates. We used the ResNet-18, normalized classifier, and Adam optimizer. We chose the 5-shot cross-domain task for visualization because the transition of validation accuracy is clearer than in other tasks. We set the learning rates as 0.01, 0.001, and 0.0001, and conducted four trials with randomly selected validation classes and support sets. Note that classification accuracy can be significantly changed by the randomly selected classes and samples. Therefore, we focused on the transition of the validation accuracy rather than the validation accuracy itself.




Figure 3: Classification accuracy for novel classes with different optimizers in the high-resolution single-domain task. We used the ResNet-18 as a feature extractor. This shows that the adaptive gradient optimizer (blue bars) achieved higher accuracy than other methods (red bars), particularly with the normalized classifier.

reduced the effect of a large domain shift in the cross-domain task by updating the entire network for novel classes.

3.7 INCREASING PERFORMANCE BY FINE-TUNING

We revealed in Section 3.5 and 3.6 that the fine-tuning method achieved high few-shot classification accuracy in many cases. In this section, we discuss the means of improving the performance of the fine-tuning method. We experimentally show that:

- Using a low learning rate for fine-tuning stabilizes the retraining process.
- Using adaptive gradient optimizers such as Adam increases the classification accuracy.
- Higher performance can be obtained by updating the entire network when a large domain shift occurs.

Low Learning Rate

A learning rate is a critical parameter in training a network; this is also true for fine-tuning for few-shot learning. Figure 2 shows that the retraining process can be stabilized by using a lower learning rate. For example, the transition of the validation accuracy was unstable when the learning rate was set as 0.01 and 0.001. However, the validation accuracy increased in a stable manner when the learning rate was set as 0.0001, which is lower than that used in the pretraining phase ($lr = 0.001$). This means that the learning rate for few-shot fine-tuning should be set low.

Adaptive Gradient Optimizer

Here, we show that optimizers affect few-shot classification performance when the network is




Figure 4: Relationship between test accuracy and updated part of the network in the 5-shot cross-domain task. We used the Adam optimizer for fine-tuning. The result for “VGG-16 Normalized BN & FC” was left empty because the network did not have a BN layer. In addition, we did not conduct experiments for the VGG-16 with the simple classifier because the loss did not decrease in the pretraining phase.

fine-tuned with only a few examples. In this experiment, we used the well-known optimizers of Adam (Kingma & Ba, 2014), Adamax (Kingma & Ba, 2014), Adadelta (Zeiler, 2012), Adagrad (Duchi et al., 2011), RMSprop (Graves, 2013), Momentum-SGD, and ASGD (Polyak & Juditsky, 1992). Of these, Adam, Adamax, Adadelta, Adagrad, and RMSprop are known as adaptive gradient methods. Figure 3 shows the classification accuracies for a 5-shot high-resolution single-domain task using the ResNet-18 with different optimizers. The results show that higher classification accuracies could be obtained by using adaptive gradient optimizers, particularly when the normalized classifier was used. Although Wilson et al. (2017) revealed that local minima obtained by the Adam optimizer lack a generalization ability, our experimental results show that this was not necessarily true for few-shot learning. Why this occurs in few-shot learning is interesting, but this is beyond the scope of this study. Therefore, we leave this interesting direction for future works.

Updating the Entire Network for Adaptation to a Large Domain Shift

Figure 4 shows the relationship between test accuracy and the updated parts of the network. This shows that updating the entire network achieves higher accuracy, particularly when the normalized classifier is used. Considering the results for the high-resolution single-domain task, when test accuracy was not further increased by updating the entire network, it could be argued that updating the entire network when a large domain shift occurs between base and novel classes is preferable.

4 CONCLUSION

In this study, we showed that in the low-resolution single-domain task, our fine-tuning method achieved higher accuracy than common few-shot learning methods in the 1-shot task and nearly the same accuracy as the state-of-the-art method in the 5-shot task. We also evaluated our method with more practical tasks, such as the high-resolution single-domain and cross-domain tasks. In both tasks, our method achieved higher accuracy than common few-shot learning methods. We then experimentally showed that: 1) a low learning rate stabilizes the retraining process, 2) adaptive gradient optimizers such as Adam improve test accuracy, and 3) updating the entire network results in higher accuracy when a large domain shift occurs. We believe that these insights into fine-tuning for few-shot learning tasks will help our community tackle this challenging task.

REFERENCES

- Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019.
- 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*, 2009.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12, 2011.

- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
- 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*, 2016.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8), 1997.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*. 2012.
- Atsuhiko Noguchi and Tatsuya Harada. Image generation from small datasets via batch statistics adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4), 1992.
- Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Abhishek Kumar, Rogerio Feris, Raja Giryes, and Alex Bronstein. Delta-encoder: an effective sample synthesis method for few-shot object recognition. In *Advances in Neural Information Processing Systems*, 2018.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 2017.
- Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- 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*, 2018.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, 2016.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, California Institute of Technology, 2011.

Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*, 2017.

Matthew D Zeiler. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.