

Proxy Network for Few Shot Learning

Bin Xiao

BUPTXIAOFENG@GMAIL.COM

Department of Computer Science, National Chiao Tung University, 1001 University Road, Hsinchu 30010, Taiwan, ROC

Chien-Liang Liu

CLLIU@MAIL.NCTU.EDU.TW

Department of Industrial Engineering and Management, National Chiao Tung University, 1001 University Road, Hsinchu 30010, Taiwan, ROC

Wen-Hoar Hsaio

BASS28.CS96G@G2.NCTU.EDU.TW

Information Management Center, National Chung-Shan Institute of Science and Technology, Taoyuan 32546, Taiwan, ROC

Editors: Sinno Jialin Pan and Masashi Sugiyama

Abstract

The use of a few examples for each class to train a predictive model that can be generalized to novel classes is a crucial and valuable research direction in artificial intelligence. This work addresses this problem by proposing a few-shot learning (FSL) algorithm called proxy network under the architecture of meta-learning. Metric-learning based approaches assume that the data points within the same class should be close, whereas the data points in the different classes should be separated as far as possible in the embedding space. We conclude that the success of metric-learning based approaches lies in the data embedding, the representative of each class, and the distance metric. In this work, we propose a simple but effective end-to-end model that directly learns proxies for class representative and distance metric from data simultaneously. We conduct experiments on CUB and mini-ImageNet datasets in 1-shot-5-way and 5-shot-5-way scenarios, and the experimental results demonstrate the superiority of our proposed method over state-of-the-art methods. Besides, we provide a detailed analysis of our proposed method.

Keywords: Few-shot Learning, Metric Learning, Proxy Networks

1. Introduction

The last decade has witnessed great success of deep learning (Krizhevsky et al., 2012; LeCun et al., 2015), because it has achieved promising results in many application domains. In image domain, convolutional neural networks (LeCun et al., 1998; Krizhevsky et al., 2012; LeCun et al., 2015; He et al., 2016) (CNNs) can benefit from their shared-weights architecture and translation invariance characteristics to learn discriminative and hierarchical features from data without human intervention. However, thousands of labeled examples for each class are required to saturate performance in classification problems, but this differs from human learning. Humans can recognize new object classes even though only few instances are available at hand. Thus, the use of a few examples for each class to train a predictive model that can be generalized to novel classes is a crucial and valuable research direction.

Many recent works proposed to use few-shot learning (FSL) and meta learning (Fei-Fei et al., 2006; Fink, 2005; Triantafillou et al., 2017; Snell et al., 2017; Finn et al., 2017) to deal with aforementioned problem. A traditional machine learning algorithm aims to learn and optimize a single task, whereas a meta-learning model is trained over a variety of learning tasks and optimized for the best performance on a distribution of tasks, including potentially unseen tasks (Weng, 2018). Notably, the problem pertaining to how to quickly learn a model for generalizing it to a new task is the core objective of meta learning, explaining why meta learning is also known as learning to learn.

Many previous FSL works (Vinyals et al., 2016; Snell et al., 2017; Sung et al., 2018) have employed the meta-learning architecture to train and evaluate the model by using various tasks. An immediate model can be obtained through learning from a variety of tasks, and the learned model can quickly adapt to a new task by using a few data examples (Finn et al., 2017). Overcoming this problem is also crucial for the development of artificial intelligence to quickly learn new things with prior knowledge as humans do. Similarly, the FSL models can use the prior knowledge that has learned from various tasks to constrain the search space of the new task (Wang and Yao, 2019).

Metric-based approaches are a class of methods that can be used to attack FSL problems, and the goal is to learn an embedding space that makes the data points in the same class to be close, whereas the data points in different classes to be separated as far as possible. We analyze the previous metric-based approaches on FSL problems and conclude that the success of metric-based approaches lies in three components: data embedding, the representative of each class, and the distance metric. We propose to use data-driven approaches to develop a novel model called proxy network (ProxyNet) under meta-learning settings. In our proposed ProxyNet, the data embeddings, and the proxies for class representatives as well as distance metric are all learned from data, and these differ from the previous metric-based approaches on FSL. The proposed model is simple to implement, but the experimental results demonstrate the superiority of our proposed method over state-of-the-art methods on 1-shot-5-way and 5-shot-5-way problems.

The contributions of this work are listed as follows. First, we analyze the previous metric-based approaches on FSL problems and conclude the key elements that lead to the success of metric-based approaches. Then, we propose a novel FSL algorithm called ProxyNet based on our analysis and the meta-learning settings. Second, we conduct experiments on two well-known benchmark datasets, and compare our proposed method with recent state-of-the-art methods. The experimental results indicate that our proposed method significantly outperforms these methods. Finally, we provide detailed analysis and investigation about our proposed method.

2. Related Work

Few-shot learning has attracted much attention in recent years as the goal is to recognize novel classes by using very few training examples (Fei-Fei et al., 2006; Vinyals et al., 2016; Snell et al., 2017; Finn et al., 2017; Sung et al., 2018). Notably, a typical CNN architecture comprises feature extractor and a classifier, and several previous works proposed to modify the classifier to deal with novel classes with a few examples. For example, Qi et al. (2018) proposed an imprinting method to directly set the weights of final layer for novel classes

from the embeddings of training exemplars. [Gidaris and Komodakis \(2018\)](#) proposed to extend CNN with an attention-based FSL weight generator and used cosine similarity between feature representations and classification weight vectors to determine prediction outcomes. [Wang and Yao \(2019\)](#) conducted a detailed survey about recent FSL works and categorized these methods into several approaches, such as metric-based, optimization-based, and augmentation-based methods.

Metric-based algorithms address the few-shot classification problem by learning to compare. [Koch et al. \(2015\)](#) combined a siamese architecture ([Chopra et al., 2005](#)) and a convolutional architecture to develop deep siamese networks for learning data representations and modeling few-shot learning as a verification task. Different from traditional deep learning models, the siamese network receives two inputs every time, and the two inputs are fed into the same network architecture with the same weights to transform the inputs into an embedding space. Once the transformation is completed, the objective is to make the data points that are in the same class to be close and the data points in different classes to be far away. [Hoffer and Ailon \(2015\)](#) extended the idea of siamese networks to devise triplet networks that accept three data points as the inputs at a time, including anchor data (reference point), positive data, and negative data. The aim of triplet loss is to bring the reference point near the positive point and away from the negative point.

Matching networks ([Vinyals et al., 2016](#)) were the first to employ the meta-learning architecture that trains and tests using the setting of K -shot- N -way tasks, in which N is the number of classes in each task, and K represents the labeled examples per class. Matching networks mapped a small number of labeled examples in the support set and unlabeled examples in the query set to embedding spaces with different embedding architectures. Cosine similarity between the embeddings of the data points in the support and query sets were used to determine the classification outcome. [Snell et al. \(2017\)](#) proposed prototypical networks by using the mean embedding or prototype as the representation for each class. The Euclidean distance between query embedding and the mean embedding for each class becomes the basis of the classification. Different from matching networks and prototypical networks that use pre-defined distance metrics to measure the similarity between the embeddings, the relation network model proposed by [Sung et al. \(2018\)](#) comprises two parts, embedding and relation modules. The purpose of embedding module was the same as those of aforementioned metric-based FSL algorithms, while the relation module used a learnable distance metric to calculate the similarity of embeddings by obtaining relation scores. Notably, the prediction of relation scores is considered a regression problem even though the targets of ground-truth are the values of 0 or 1.

Optimization-based methods constrain the learning algorithms by providing a good initialization of the model parameters, or learning the update rules. Model agnostic meta-learning (MAML) ([Finn et al., 2017](#)) is a meta-learning framework that learns a network initialization that can quickly adapt to new tasks. [Nichol et al. \(2018\)](#) introduced a variant of first order MAML called Reptile and explored the potential of meta-learning algorithms based on first-order gradient information. Variants of gradient-based optimization algorithms are standard optimization algorithms to train deep neural networks. [Ravi and Larochelle \(2016\)](#) discovered that the update rule for the variants of gradient descent resembles the update rule for the cell state in a long short-term memory (LSTM) ([Hochreiter](#)

and Schmidhuber, 1997), inspiring them to propose an LSTM-based meta-learner optimizer that was trained to optimize a NN classifier.

Augmentation-based methods cope with FSL by hallucinating new examples to augment datasets for the novel classes. Intuitively, this approach relies on the quality of the generated examples. Hariharan and Girshick (2017) considered to use analogy relationship to generate new examples. Given quadruplets of features taken from training tasks, the aim was to train a generator that can output a hallucinated feature vector by transferring modes of variation from the base classes. Subsequently, Wang et al. (2018) proposed a method that jointly trained a meta-learner and a hallucinator, which mapped real seed examples and noise vectors to hallucinated examples.

3. Proposed Method

This section briefly introduces FSL problems in the meta-learning framework. Next, we analyze the components involved in the previous metric-based methods on FSL and present an overall architecture as well as these key components for metric-based methods. Based on this architecture, we further present our proposed model, proxy networks (ProxyNet).

3.1. Problem Definition

In this study, we follow meta-learning settings to design the model and conduct experiments. First, we separate the process into meta-training and meta-testing stages. Notably, the datasets in meta-training and meta-testing have disjoint classes as the purpose is to evaluate whether the model can make accurate predictions on novel classes. Second, the training and evaluation are conducted with a collection of K -shot- N -way tasks in episodes. Conventionally, K is less than 20 in the setting of FSL, and the goal is to assess whether the model can quickly adapt to new tasks with few labeled examples.

In meta-training stage, each task comprises a support set, $S = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N \cdot K}$, where \mathbf{x}_i is the input example and $y_i \in \{1, \dots, N\}$ is the corresponding label, and a query set, $Q = \{(\mathbf{x}_j, y_j)\}_{j=1}^{N \cdot T}$, where T is the number of query instances per class in each meta-training episode, \mathbf{x}_j is the query example and $y_j \in \{1, \dots, N\}$ is the corresponding label. Additionally, $S \cap Q = \emptyset$, but S and Q share the same label space. The algorithm has to determine which of the support set class c the query samples q belongs to based on the class assignment probability $P_\phi(y = c|q)$, in which ϕ is model parameter. The settings of meta-testing resemble those of the meta-training stage except that the query set does not comprise label y .

3.2. Metric-based FSL

This work summarized previous works that used metric-based approaches to deal with FSL problems and Fig. 1 shows the overall architecture for these methods. A metric-based FSL model typically comprises three key components, including embedding functions (f_θ is for support set and g_θ is for query set), class representative function (h_θ), and the distance metric function (d_θ) that is used to determine the similarity of support and query embeddings. We aim to use non-parametric model to determine the class assignment as the work

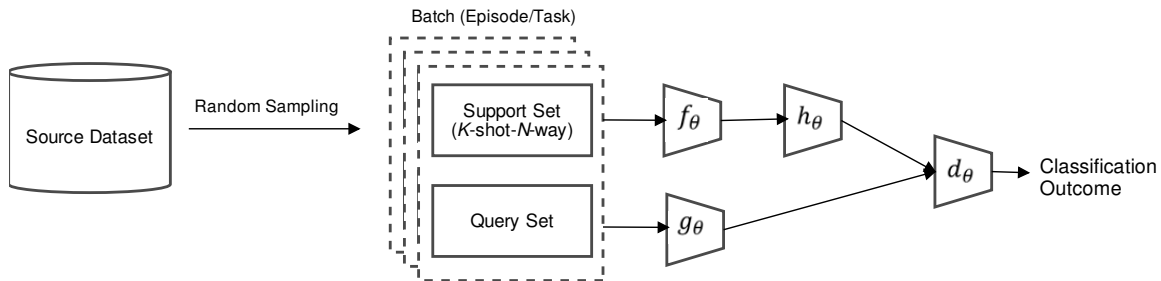


Figure 1: Metric-based FSL Architecture.

conducted by [Chen et al. \(2019\)](#) has empirically shown that a simple distance-based classifier could achieve competitive performance with the state-of-the-art meta-learning methods. Thus, [Fig. 1](#) uses a distance metric function (d_θ) to perform classification without considering parametric functions.

It is apparent that embedding functions are crucial to the subsequent classification results. A common approach is to use a simple network such as Conv-4 as the embedding function, but a deep neural network such as ResNet-32 or a fully conditional embedding network such as bi-LSTM could be the alternative. Moreover, f_θ and g_θ can be different as used in matching network ([Vinyals et al., 2016](#)), and be identical as used in prototypical networks ([Snell et al., 2017](#)) and relation networks ([Sung et al., 2018](#)).

In a K -shot- N -way task, each class comprises several embedding vectors when $K > 1$, and h_θ accounts for the generation of the class representative for each class. Given the embedding vectors for each class, the h_θ in prototypical network is an average operator and outputs the class mean of these embedding vectors. As for relation network, the h_θ is a sum operator. On the other hand, matching network focuses on 1-shot learning, namely, $K = 1$. In this case, the h_θ is an identity operator that directly outputs the input embedding vector.

The final step for the metric-based FSL is to perform classification, and d_θ can be a classifier or distance metric. Parametric models need to be optimized to treat new examples. In contrast, non-parametric models such as nearest neighbor method can perform classification only depending on the chosen metric, explaining why many metric-based methods used nearest neighbor method with appropriate distance metric to perform classification. Among these methods, matching networks and prototypical networks used pre-defined metrics, in which the former used cosine similarity, whereas the latter used Euclidean distance. Relation networks used a different approach by learning a distance metric from the relation module.

3.3. Proxy Networks

Based on the overall architecture introduced above, this work proposes a metric-based FSL method called ProxyNet. We propose to use the same embedding function to map support and query sets to an embedding space, namely, $f_\theta = g_\theta$. We follow the experimental protocol used in the work conducted by [Chen et al. \(2019\)](#) to use a Conv-4 network to conduct experiments and compared ProxyNet with other alternatives. Notably, different

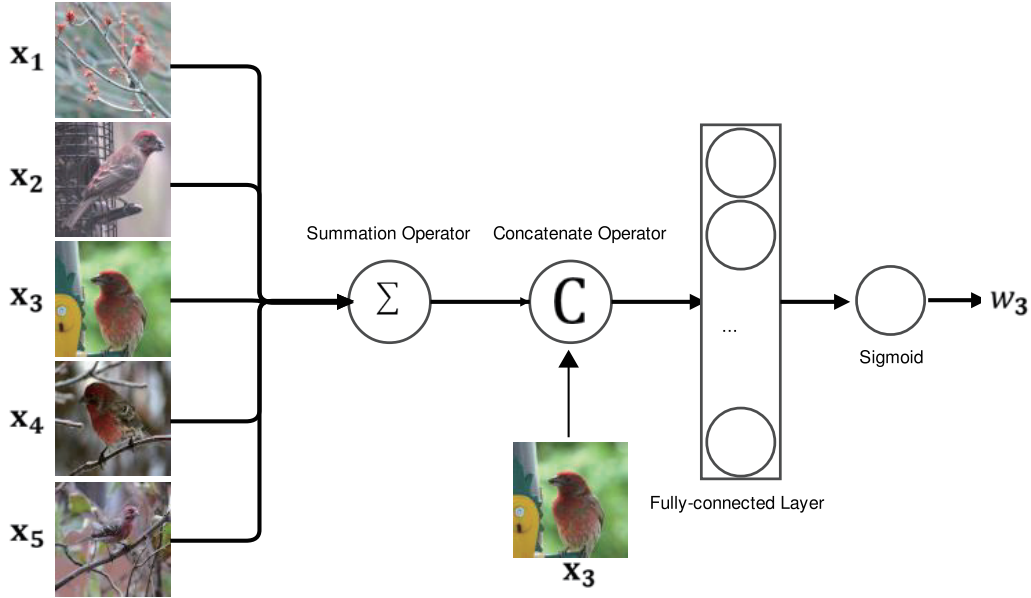


Figure 2: The architecture for the estimation of data weight in class proxy.

embedding functions can be used in our proposed model and we conduct experiments to evaluate the impact of different embedding networks on the performances. Apart from the embedding functions, we propose to directly learn class representatives from data instead of using pre-defined operators such as average or sum to obtain the representatives. Given a class with K points, $\mathbf{x}_1, \dots, \mathbf{x}_K$, the center of the K points is a commonly used point to represent the entire class. However, this approach will tend to be biased by outliers, so we propose a learning-based approach to learn the representative for each class. To simplify the notations for the following steps of our algorithm, the notations for the input images and their corresponding embeddings use the same notation \mathbf{x} . Given a K -shot- N -way task, there exists $\mathbf{x}_1^n, \dots, \mathbf{x}_K^n$ examples for n th class, and we propose to use a function h_θ to learn the class representative of n th class μ_n , namely, $\mu_n = h_\theta(\mathbf{x}_1^n, \dots, \mathbf{x}_K^n)$, where $n \in \{1, \dots, N\}$. Following the same setting, one can obtain all the class representatives, μ_1, \dots, μ_N , from h_θ . We define the function as a weighted sum of the input points, namely, $h_\theta(\mathbf{x}_1^n, \dots, \mathbf{x}_K^n) = w_1^n \mathbf{x}_1^n + w_2^n \mathbf{x}_2^n + \dots + w_K^n \mathbf{x}_K^n$, where w_k^n ($k \in \{1, 2, \dots, K\}$) is the weight of the k th support point in the n th class. We propose to use a deep neural network as shown in Fig. 2 to learn the weight w_k^n , in which the inputs of the network comprises \mathbf{x}_k^n and the summation of the data points for the n th class, say \mathbf{s}^n . The goal is to learn the similarity score between \mathbf{x}_k^n and \mathbf{s}^n , since \mathbf{s}^n can be considered an approximation to the global trend of the data points in the n th class. Intuitively, the data point that is more close to the global trend of the class, more weight should be given to it. One advantage of such a design is to alleviate the influence of outliers by giving small weights to the outliers; in contrast, many pre-defined operators, such as sum or mean, are easily influenced by the outliers.

Finally, metric-based approaches require to calculate the similarity between query data and class proxies, so that the class assignment can be obtained. Similarly, we propose to use a distance network (Sung et al., 2018) to learn the similarity or distance between query data

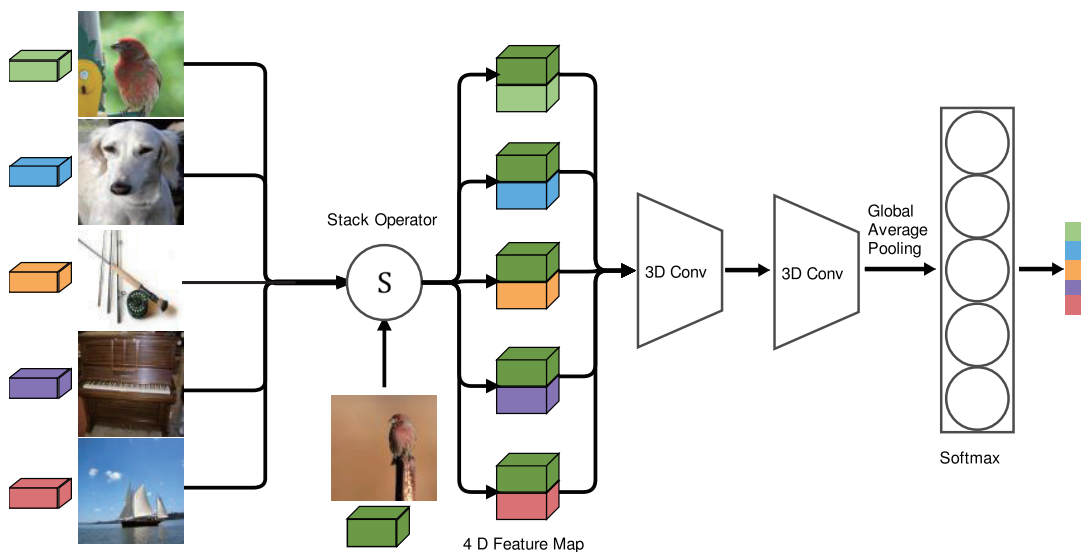


Figure 3: Classification module based on 3D Convolution.

and each class proxy. Fig. 3 shows the architecture of our classification module. Given the class proxies and a query embedding, we stack the query embedding and each class proxy embedding on a new dimension, so that the dimension is four after the stacking operator. Subsequently, we use a 3D convolutional network and global average pooling to obtain the probabilities of the class assignment.

As shown in Fig. 3, the 3D convolution network comprises two 3D convolutional blocks, each of which comprises 3D convolution, 3D batch normalization, and ReLU activation function. The aim is to learn the relation between channels and the features in each channel. Following the relation learning, a global average pooling layer is used to compute the average value for each feature map and supplies it to a softmax layer. Notably, we use a global average pooling layer to replace fully connected layers as used in (Lin et al., 2013), since it can decrease model complexity and reduce overfitting. Moreover, we list the number of trainable parameters for the proposed ProxyNet in the experimental section, showing the superiority of ProxyNet over state-of-the-art methods while requiring fewer trainable parameters. The output of the softmax layer comprises the probabilities of class assignments, and then the cross entropy loss function is used to train our proposed model. It is worth mentioning that the proposed model can be trained in an end-to-end manner.

4. Experiments

This section introduces the datasets that used in the experiments, the experimental settings, and the experimental results on FSL problem.

4.1. Datasets

We evaluate the proposed model on two datasets, CUB (Wah et al., 2011) and mini-ImageNet (Vinyals et al., 2016). Notably, these two datasets are commonly used by other

previous FSL works. The CUB dataset comprises 100 classes for training, 50 classes for validation, and 50 classes for testing. This setting is the same as the evaluation protocol of (Hilliard et al., 2018). The mini-ImageNet dataset consists of 64 classes for training, 16 classes for validation, and 20 classes for testing. Each class has 600 images, each of which is of the size 84×84 . Accuracy is used as the evaluation metric in the experiments.

4.2. Experimental Settings

We follow the experiment settings listed in (Chen et al., 2019) to conduct experiments on mini-ImageNet and CUB datasets. First, all the methods use Conv-4 as the embedding network, which is consisted of 4 convolution layers and each convolution layer has 64 filters of size 3×3 , since the purpose is to focus on the metric-based models on FSL problem rather than the effectiveness of the embedding network. Second, we evaluate all the methods on two scenarios, 1-shot-5-way and 5-shot-5-way, as they are mostly used evaluation scenarios in FSL research. Third, we randomly generate 600 test tasks from testing set, and the number of query images in each class is 15. Thus, for a 5-way-5-shot classification problem, each task comprises 75 query images as each class contains 15 images. We use the validation set to determine the best models in the experiments, in which we randomly generate 600 tasks from the validation set. The optimizer in our proposed model is SGD with an initial learning rate of 0.1 and the learning rate is adjusted by a reduce-on-plateau (PytorchDevTeam, 2020) method during training. The training stage comprises 300 epochs for the CUB dataset and 600 epochs for the mini-ImageNet dataset, each of which has 100 episodes. We use the same preprocessing methods as Chen et al. (2019) did in their work. More specially, all the images in both datasets are firstly resized to 92×92 and random cropped to 84×84 , and then further processed by color jitters and random horizontal flip. The mean accuracy of generated tasks and their 95% confidence intervals are presented as the experimental results. We have released the source code of our model on github ¹, and detailed hyper-parameters as well as settings are available in the source code.

4.3. Experimental Results

We conduct experiments on two datasets, and compare the proposed ProxyNet with several alternatives. The proposed method is a metric-based method, so all comparison methods are metric-based methods, including MatchingNet, ProtoNet, RelationNet, category traversal module (CTM) (Li et al., 2019a), deep nearest neighbor neural network (DN4) (Li et al., 2019b), and baseline++. The introduction for the first three comparison methods can refer to the previous section. Notably, these comparison methods use non-parametric approaches to determine the class assignment, and they can be modeled by the proposed framework as listed in Fig. 1.

CTM (Li et al., 2019a) leveraged a concentrator and a projector to learn a mask which can be applied to all the support and query samples to obtain discriminative features. Instead of learning class representatives and comparing the query samples with class representatives, CTM compares all the support and query sample pairs and uses the average distance between a query sample and all the support samples within a class as the distance

1. ProxyNet: https://github.com/buptxiaofeng/proxynet_fsl

to that class. Notably, we implement the three components of CTM, including concentrator, projector, and reshaper, with three residual blocks for a fair comparison, and the purpose is to make the number of trainable parameters closer to other models. Besides, the implementation for CTM uses Conv-4 as the embedding network, and uses two convolution layers as the distance metric as used by RelationNet.

DN4 (Li et al., 2019b) treated the feature maps generated by the embedding network of each sample as a series of local descriptor and proposed a Image-to-Class module to calculate the image-to-class similarity. Given a query sample, the local descriptors are firstly used to find the k nearest neighbors in a class. Subsequently, the sum of similarities between the local descriptors of the query sample and the k nearest neighbors is used to measure the image-to-class similarity between the query sample and the class. Our implementation for DN4 is based on the settings used by other comparison methods. More specifically, we used Conv-4 as the embedding network and used random crop, color jitters and random horizontal flip as the augmentation methods during training. Notably, we used 1-nearest neighbors on the CUB dataset and 3-nearest neighbors on the mini-ImageNet dataset as these settings were used by the original DN4.

The baseline++ was proposed by Chen et al. (2019), and was developed based on transfer learning of network pre-training and fine-tuning. In training stage, the tasks in the meta-training dataset were used to train the underlying embedding network and a classifier of parameter $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_C]$, where C is the number of classes or ways in the task. The class assignment is based on the cosine similarity between the query set and \mathbf{w}_c ($1 \leq c \leq C$).

The parameters $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_C]$ resembles the C prototypes used by ProtoNet, but differs in that \mathbf{W} is learned from data. In fine-tuning stage, they used the support sets in the meta-testing dataset to fine-tune the classifier. Among the three components listed in Fig. 1, baseline++ trains the underlying embedding network and the class representatives, but uses a pre-defined distance function, cosine.

Table 1 shows the experimental results. Notably, the experimental results for baseline++, MatchingNet, ProtoNet, and RelationNet are obtained from the work conducted by Chen et al. (2019), and we follow their experimental settings for a fair comparison, including data augmentation methods and the embedding networks. More specially, we resize the images to the size of 92×92 followed by cropping to the size of 84×84 . Subsequently, all the images are further processed by color jitters and random horizontal flip. The experimental results listed in Table 1 indicate that the proposed method significantly outperforms baseline++, MatchingNet, ProtoNet, and RelationNet. Meanwhile, the proposed method also generally achieves better results than CTM and DN4, but the performance differences between CTM and ProxyNet are insignificant in the two problems of the CUB dataset. It is worth mentioning that the concentrator of CTM can be treated as part of the embedding network which makes the embedding network deeper than the others. The other components of CTM also employed deeper networks, resulting in more trainable parameters than the proposed method. We list the numbers of trainable parameters for all the models in Table 2, which indicates that the number of trainable parameters for CTM is much more than that of ProxyNet. We provide detailed analysis regarding the impact of deeper embedding networks on the performance of ProxyNet in the following section. We summarize the three components and the learning type of the metric-based models and the proposed

Table 1: Experimental results on 1-shot-5-way and 1-shot-5-way tasks.

Method	CUB		mini-ImageNet	
	1-shot	5-shot	1-shot	5-shot
ProxyNet	67.52 \pm 0.97	82.85 \pm 0.60	52.95 \pm 0.76	71.02 \pm 0.62
CTM (Li et al., 2019a)	68.61 \pm 0.95	82.72 \pm 0.58	50.80 \pm 0.84	67.79 \pm 0.67
DN4 (Li et al., 2019b)	52.01 \pm 0.86	76.70 \pm 0.64	47.32 \pm 0.80	69.22 \pm 0.66
Baseline++ (Chen et al., 2019)	60.53 \pm 0.83	79.34 \pm 0.61	48.24 \pm 0.75	66.43 \pm 0.63
MatchingNet (Vinyals et al., 2016)	60.52 \pm 0.88	75.29 \pm 0.75	48.14 \pm 0.78	63.48 \pm 0.66
ProtoNet (Snell et al., 2017)	50.46 \pm 0.88	76.39 \pm 0.64	44.42 \pm 0.84	64.24 \pm 0.72
RelationNet (Sung et al., 2018)	62.34 \pm 0.94	77.84 \pm 0.68	49.31 \pm 0.85	66.60 \pm 0.69

Table 2: The number of trainable parameters.

ProxyNet	CTM	DN4	Baseline++	MatchingNet	ProtoNet	RelationNet
165, 171	307, 921	112, 832	433, 288	138, 176	113, 600	229, 201

method in Table 3, in which the transfer learning in the table means that Baseline++ follows the typical paradigm of transfer learning to train the model.

Table 3: Summary of metric-based learning models and the proposed model.

Method	Embedding Network	Class Proxy	Distance Metric	Type
ProxyNet	ConvNet	trainable proxy	3D-Conv and global pooling	meta learning
DN4	ConvNet	No Proxy	Image-to-Class module(Cosine Distance)	meta learning
Baseline++	ConvNet	No Proxy	Cosine Distance	transfer learning
MatchingNet	LSTM	No Proxy	Cosine Distance	meta learning
ProtoNet	ConvNet	Mean of support samples	Euclidean Distance	meta learning
RelationNet	ConvNet	Sum of support samples	2D-Conv and fully connected layer	meta learning

5. Analysis and Discussion

The experimental results in the previous section show that the proposed method can benefit from the learnable proxies for class representatives and learnable distance function. Moreover, the previous experiments only use simple CNNs as the embedding networks, deeper CNNs can be also the alternatives. Thus, we provide detailed analysis and discussion regarding the underlying embedding networks, proxy for class representative, and proxy for distance function in this section.

5.1. Embedding Network

In the previous experiments, the embedding network is a simple embedding network, Conv-4. We investigate whether the ProxyNet could benefit from a deeper CNN, so we conduct experiments by applying ProxyNet with different embedding networks to the two datasets that have undergone the data augmentation process as mentioned in the previous section. Table 4 shows the experimental results, indicating that ProxyNet can normally benefits from deeper networks to obtain better performances. Although this work focuses on the design

Table 4: Experimental results on 5-way tasks using different embedding networks with augmentation methods.

Embedding Network	CUB		mini-ImageNet	
	1-shot	5-shot	1-shot	5-shot
Conv-4	67.52 ± 0.97	82.85 ± 0.60	52.95 ± 0.76	70.35 ± 0.63
Conv-6	68.16 ± 0.93	83.57 ± 0.58	52.18 ± 0.82	69.91 ± 0.62
ResNet-10	76.79 ± 0.84	88.02 ± 0.52	58.16 ± 0.87	75.27 ± 0.65
ResNet-18	76.72 ± 0.90	88.63 ± 0.49	57.88 ± 0.87	75.23 ± 0.66
ResNet-34	77.70 ± 0.86	87.05 ± 0.52	58.54 ± 0.85	75.90 ± 0.61

of networks for class representatives and distance function, it is apparent that the design of embedding network is also crucial to the model performance. Notably, data augmentation is a commonly used technique that can enable practitioners to significantly increase the size and diversity of the data, making it possible to train a deeper network under the FSL settings.

Besides, we conduct experiments to investigate whether ProxyNet can benefit from a deeper CNN when no data augmentation is applied to the two datasets. The purpose is to analyze the relationship between the choice of embedding networks and the available data sizes. The experimental results are listed in Table 5.

Table 5: Experimental results on 5-way tasks using different embedding networks without augmentation methods.

Embedding Network	CUB (No data augmentation)		mini-ImageNet (No data augmentation)	
	1-shot	5-shot	1-shot	5-shot
Conv-4	63.95 ± 0.95	79.41 ± 0.64	52.10 ± 0.89	71.02 ± 0.62
Conv-6	63.22 ± 0.91	77.24 ± 0.70	52.29 ± 0.82	69.52 ± 0.66
ResNet-10	62.49 ± 0.95	76.15 ± 0.75	53.27 ± 0.82	69.10 ± 0.67
ResNet-18	63.47 ± 0.95	72.50 ± 0.78	53.64 ± 0.82	68.86 ± 0.66
ResNet-34	62.68 ± 1.01	71.83 ± 0.81	52.69 ± 0.81	68.47 ± 0.66

The experimental results in Table 5 show an opposite trend of performances against network depth, meaning that the ProxyNet with simple CNN can normally yield better performance than that with deep CNNs. We conjecture that ProxyNet can not always benefit from a deep embedding network owing to the overfitting problem. The ProxyNet is a flexible design as it learns embedding network, class proxy, and distance metric proxy from data. This brings the model more capacity to cope with FSL problem, but also increases the model parameters. Thus, ProxyNet is prone to overfitting problem when the available samples are insufficient.

Moreover, the results in Table 5 show that it seems like that overfitting problem is more serious in CUB dataset than mini-ImageNet. The CUB dataset is designed for fine-grained image classification problem, meaning that the domain difference between base classes and

novel classes is relative smaller than that of mini-ImageNet. Thus, FSL problem on CUB dataset is easier than that on mini-ImageNet. Besides, it is apparent that 5-shot-5-way problem is easier than 1-shot-5 way problem. Thus, it is expected that 5-shot-5-way problem on CUB dataset is the easiest problem, whereas 1-shot-5-way problem on mini-ImageNet is the most difficult problem among the problems listed in Table 5. Based on the problem difficulties and the experimental results presented in Table 5, it seems like that the ProxyNet with a deep embedding network is prone to overfitting when coping with easy problems.

We conduct experiments to empirically demonstrate that ProxyNet can benefit from data augmentation to dramatically improve performance as compared with that does not use data augmentation. Moreover, the findings show that ProxyNet should provide sufficient training examples when using a deep CNN as the embedding network; otherwise, the model may suffer from overfitting problem. In the setting of FSL, data augmentation is an effective approach to increase the sample size and avoid the overfitting problem.

5.2. Class Proxy

When the underlying distance metric is Euclidean distance, a typical approach to obtain a class representative is to apply the average operator to the class members, meaning that the center point can be the representative as used by ProtoNet. The sum operator is similar to average operator by summing the class members as used by RelationNet. Different from these two operators, we propose to use weighted average of the class members to obtain the class representative, but the weights are learned from a neural network.

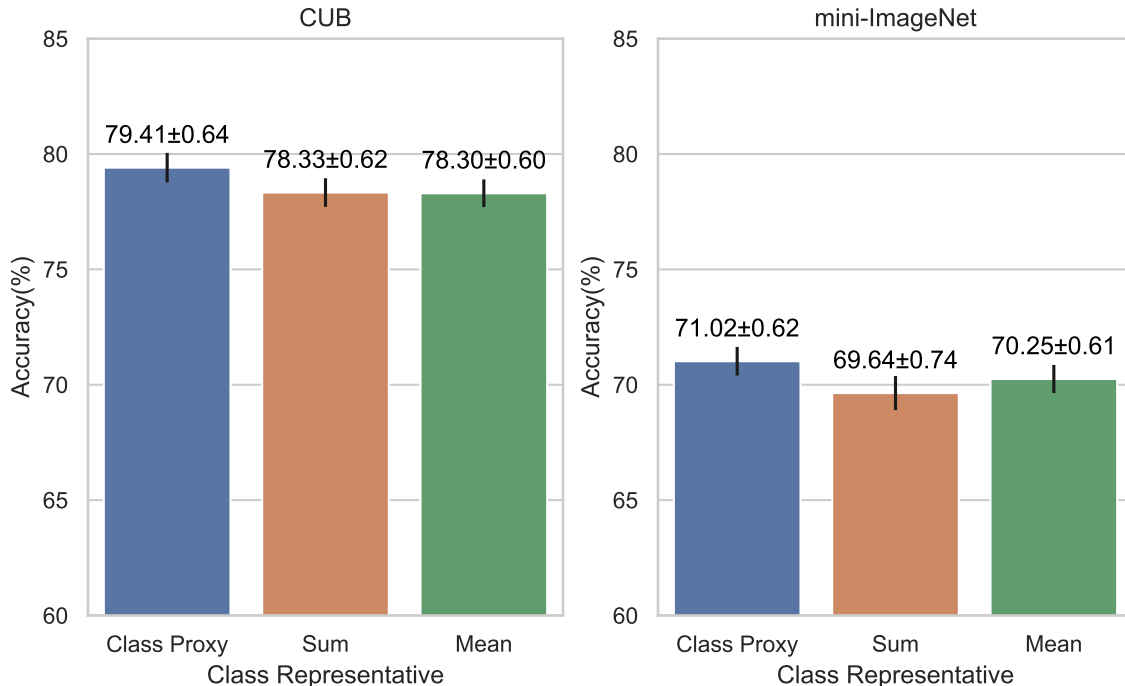


Figure 4: Performance comparison with different approaches to obtain class representatives on 5-way-5-shot task.

To investigate the effectiveness of the proposed class representative proxy, we use our proposed method with mean and sum operators to obtain class representative as two comparison methods while fixing other settings. Fig. 4 shows the experimental results, indicating that the proposed class proxy yields better performance than the other two comparison methods. Aggarwal et al. (2001) has viewed the dimensionality curse from the point of view of the distance metrics and examined the behavior of the commonly used L_k norm. They have shown that the problem of meaningfulness in high dimensionality is sensitive to the value of k and the problem worsens faster for higher values of k . Thus, using Euclidean distance (L_2 norm) to search for the nearest neighbor is not always meaningful for high dimensional data. In contrast, using a network to directly learn class proxy is more preferable than using a pre-defined operator to obtain class representative.

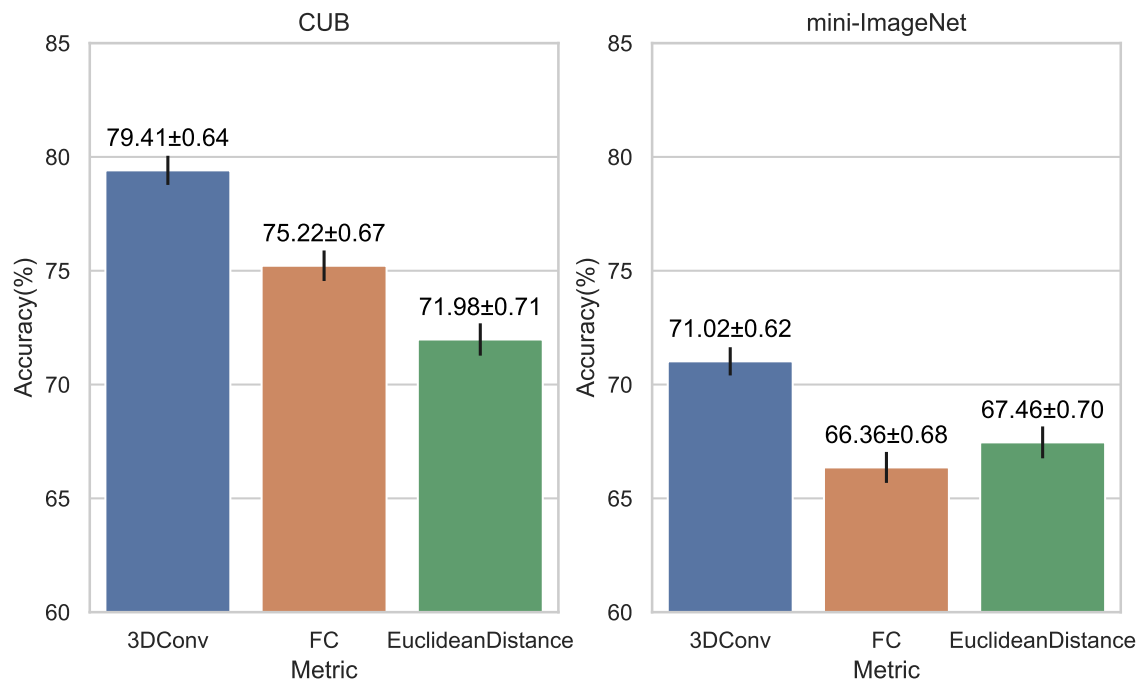


Figure 5: Performance comparison using different metrics on 5-way-5-shot task.

5.3. Distance Metric Proxy

The final step of metric-based methods is to determine the class assignment of the query embedding. Given support and query embeddings, the most commonly used approach is to measure the distances between query embedding and all the support embeddings. For example, ProtoNet uses Euclidean distance, while baseline++ uses cosine similarity. Instead of using fixed distance functions, RelationNet uses a relation module to learn the relation score from data, in which the network composed of a convolution layer, a max-pooling layer and a fully connected layer (FC in Fig. 5). This work also proposes to learn a distance function from data, but we propose to use a 3D convolutional network (3DConv in Fig. 5)

to learn the relationship from data and use this kind of relationship as the proxy of distance metric.

Fig. 5 shows the experimental results on CUB and mini-ImageNet datasets, which indicate that the proposed ProxyNet yields the best performance among these methods. The ProxyNet uses a squeeze-and-excitation (Hu et al., 2018) block, a 3D convolution, and a global average pooling layer to learn the relation scores between query embedding and support embeddings while retaining the spatial property. The model can benefit such a design by keeping the spatial feature and the relation of channels.

This section provides detailed analysis and discussion about ProxyNet, and conduct extensive experiments to empirically show that each component of our proposed model is crucial to the model architecture. Moreover, using a data-driven approach to learn the embeddings, class proxy, and distance metric proxy can directly learn the three components from data without human intervention, giving a base to develop a flexible and generalized model.

6. Conclusions

This work analyzes the previous metric-learning based model and summarizes three important components in a metric-learning based model. Based on our analysis, we propose a metric-learning based model called ProxyNet to cope with FSL problem. Central to the proposed model is to use data-driven approach to consider embeddings, class proxy, and distance metric proxy as three learning problems, making it possible to learn these components from data without human intervention. The experimental results indicate that the proposed method can yield promising results on CUB and mini-ImageNet datasets in 1-shot-5-way and 5-shot-5-way scenarios. Detailed analysis and discussion about ProxyNet are also provided in this work. The future work is to cope with FSL problem that comprises cross-domain datasets.

Acknowledgment

This work was supported in part by Ministry of Science and Technology, Taiwan, under Grant no. MOST 107-2221-E-009-109-MY2 and MOST 109-2628-E-009-009-MY3. We are grateful to the National Center for High-performance Computing for computer time and facilities.

References

- Charu C Aggarwal, Alexander Hinneburg, and Daniel A Keim. On the surprising behavior of distance metrics in high dimensional space. In *International conference on database theory*, pages 420–434. Springer, 2001.
- 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.
- Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *null*, pages 539–546. IEEE, 2005.

- Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006.
- Michael Fink. Object classification from a single example utilizing class relevance metrics. In *Advances in neural information processing systems*, pages 449–456, 2005.
- 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-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4367–4375, 2018.
- Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3018–3027, 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*, pages 770–778, 2016.
- Nathan Hilliard, Lawrence Phillips, Scott Howland, Artëm Yankov, Courtney D Corley, and Nathan O Hodas. Few-shot learning with metric-agnostic conditional embeddings. *arXiv preprint arXiv:1802.04376*, 2018.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Elad Hoffer and Nir Ailon. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*, pages 84–92. Springer, 2015.
- 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.
- Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML Deep Learning Workshop*, volume 2, 2015.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- Hongyang Li, David Eigen, Samuel Dodge, Matthew Zeiler, and Xiaogang Wang. Finding task-relevant features for few-shot learning by category traversal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–10, 2019a.

- Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo. Revisiting local descriptor based image-to-class measure for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7260–7268, 2019b.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
- Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- PytorchDevTeam. torch.optim-pytorch master document, 2020. URL https://pytorch.org/docs/stable/optim.html#torch.optim.lr_scheduler.ReduceLRonPlateau.
- 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*, pages 5822–5830, 2018.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- 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.
- Eleni Triantafillou, Richard Zemel, and Raquel Urtasun. Few-shot learning through an information retrieval lens. In *Advances in Neural Information Processing Systems*, pages 2255–2265, 2017.
- 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.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- Yaqing Wang and Quanming Yao. Few-shot learning: A survey. *arXiv preprint arXiv:1904.05046*, 2019.
- 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*, pages 7278–7286, 2018.
- Lilian Weng. Meta-learning: Learning to learn fast. *lilianweng.github.io/lil-log*, 2018. URL <http://lilianweng.github.io/lil-log/2018/11/29/meta-learning.html>.