

Prototype Rectification for Few-Shot Learning

Jinlu Liu, Liang Song, and Yongqiang Qin*

AInnovation Technology Co., Ltd.

liujinlu, songliang, qinyongqiang@ainnovation.com

Abstract. Few-shot learning requires to recognize novel classes with scarce labeled data. Prototypical network is useful in existing researches, however, training on narrow-size distribution of scarce data usually tends to get biased prototypes. In this paper, we figure out two key influencing factors of the process: *the intra-class bias* and *the cross-class bias*. We then propose a simple yet effective approach for prototype rectification in transductive setting. The approach utilizes label propagation to diminish the intra-class bias and feature shifting to diminish the cross-class bias. We also conduct theoretical analysis to derive its rationality as well as the lower bound of the performance. Effectiveness is shown on three few-shot benchmarks. Notably, our approach achieves state-of-the-art performance on both miniImageNet (70.31% on 1-shot and 81.89% on 5-shot) and tieredImageNet (78.74% on 1-shot and 86.92% on 5-shot).

Keywords: Few-Shot Learning · Prototype Rectification · Intra-Class Bias · Cross-Class Bias

1 Introduction

Many deep learning based methods have achieved significant performance on object recognition tasks with abundant labeled data provided [12, 27, 9]. However, these methods generally perform unsatisfactorily if the labeled data is scarce. To reduce the dependency of data annotation, more researchers make efforts to develop powerful methods to learn new concepts from very few samples, which is so-called Few-Shot Learning (FSL) [17, 5, 33]. In FSL, we aim to learn prior knowledge on base classes with large amounts of labeled data and utilize the knowledge to recognize few-shot classes with scarce labeled data. It is usually formed as N -way K -shot few-shot tasks where each task consists of N few-shot classes with K labeled samples per class (the support set) and some unlabeled samples (the query set) for test.

Classifying test samples by matching them to the nearest class prototype [30] is a common practice in FSL. It is supposed that an expected prototype has the minimal distance to all samples within the same class. However, the prototypes they get are always biased due to the data scarcity in few-shot scenarios. The internal factors that restrict the representation ability of the prototypes should be identified for performance improvement. Hence, we figure out the bias in

* Corresponding author.

prototype computation and accordingly propose the diminishing methods for rectification.

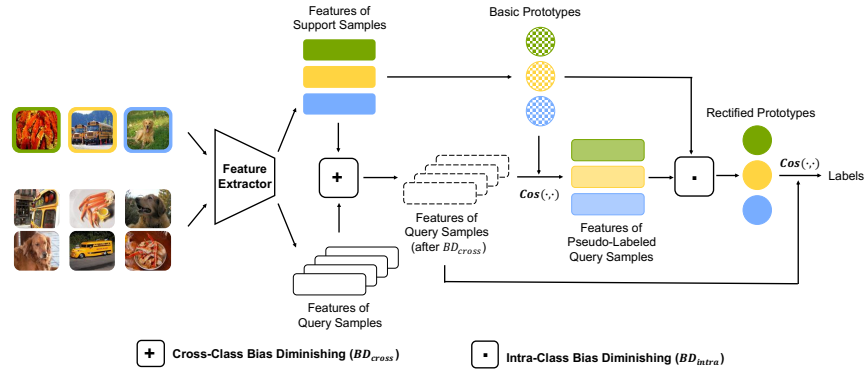


Fig. 1. Framework of our proposed method for prototype rectification. *The cross-class bias diminishing module* reduces the bias between the support set and the query set while *the intra-class bias diminishing module* reduces the bias between the actually computed prototypes and the expected prototypes.

In this paper, we target to find the expected prototypes which have the maximum cosine similarity to all data points within the same class. The cosine similarity based prototypical network (CSPN) is firstly proposed to extract discriminative features and compute basic prototypes from the limited samples. In CSPN, we firstly train a feature extractor with a cosine classifier on the base classes. The cosine classifier has a strong capability of driving the feature extractor to learn discriminative features. It learns an embedding space where features belonging to the same class cluster more tightly. At the inference stage, we use class means as the basic prototypes of few-shot classes. Classification can be directly performed by nearest prototype matching based on cosine similarity.

Since the basic prototypes are biased due to data scarcity, we import a bias diminishing module into the network for prototype rectification, which is called BD-CSPN in this paper. We figure out two key factors: the intra-class bias and the cross-class bias, which influence the representativeness of class prototypes. The approach to reduce the bias is accordingly proposed as shown in Fig. 1. The *intra-class bias* refers to the distance between the expectedly unbiased prototype and the prototype actually computed from the available data. To reduce it, we adopt the pseudo-labeling strategy to add unlabeled samples with high prediction confidence into the support set in transductive setting. Considering that some of the pseudo-labeled samples are possibly misclassified, we use the weighted sum as the modified prototypes instead of simple averaging. It avoids bringing larger bias into prototype computation. The *cross-class bias* refers to the distance between the representative vectors of training and test datasets,

which are commonly represented as the mean vectors. We reduce it by importing a shifting term ξ to the query samples, driving them to distribute closely to the support samples.

To verify the rationality of our bias diminishing method, we give the theoretical analysis in Section 4. The derivation of the expected performance of cosine-similarity based prototypical network is firstly given. It shows that the lower bound of the expected accuracy is positively correlated with the number of samples. *We demonstrate the effectiveness and simplicity of our pseudo-labeling strategy in raising the lower bound, which leads to significant improvement as shown in experiments.* Then we give the derivation of shifting term ξ in cross-class bias diminishing. In conclusion, we argue that our method is simpler yet more efficient than many complicated few-shot learning methods. Also, it is mathematically rigorous with the theoretical analysis.

Our contributions are summarized as:

- 1) We figure out the internal factors: the intra-class bias and the cross-class bias which restrict the representational ability of class prototypes in few-shot learning.
- 2) We propose the bias diminishing module for prototype rectification, which is mainly conducted by pseudo-labeling and feature shifting. It is conceptually simple but practically effective to improve the performance.
- 3) To verify the rationality of the intra-class bias diminishing method, we theoretically analyze the correlation between the number of sample and the lower bound of the expected performance. Furthermore, we give the derivation of the shifting term in cross-class bias diminishing.
- 4) We conduct extensive experiments on three popular few-shot benchmarks and achieve the state-of-the-art performance. The experiment results demonstrate that our proposed bias diminishing module can bring in significant improvement by a large margin.

2 Related Works

Few-Shot Learning Few-shot learning methods can be divided into two groups: *gradient based methods* and *metric learning based methods*. Gradient based methods focus on fast adapting model parameters to new tasks through gradient descent [2, 6, 15, 20, 28, 13]. Typical methods such as MAML [6] and Reptile [20] aim to learn a good parameter initialization that enables the model easy to fine-tune. In this section, we focus on metric learning based methods which are more closely to our approach. Metric learning based methods learn an informative metric to indicate the similarity relationship in the embedding space [33, 30, 31, 1]. Relation network [31] learns a distance metric to construct the relation of samples within an episode. The unlabeled samples thus can be classified according to the computed relation scores. Prototypical Networks (PN) [30] views the mean feature as the class prototype and assigns the points to the nearest class prototype based on Euclidean distance in the embedding space. It is indicated in [13] that PN shows limited performance in the high-dimensional embedding

space. In some recent works, models trained with a cosine-similarity based classifier are more effective in learning discriminative features [8, 3]. In this paper, we use cosine classifier to learn a discriminative embedding space and compute the cosine distance to the class prototype (mean) for classification. The prototype computed in the discriminative feature space is more robust to represent a class.

According to the test setting, FSL can be divided into two branches: *inductive few-shot learning* and *transductive few-shot learning*. The former predicts the test samples one by one while the latter predicts the test samples as a whole. Early proven in [10, 36], transductive inference outperforms inductive inference especially when training data is scarce. Some literatures recently attack few-shot learning problem in transductive setting. In [20], the shared information between test samples via normalization is used to improve classification accuracy. Different from [20], TPN [16] adopts transductive inference to alleviate low-data problem in few-shot learning. It constructs a graph using the union of the support set and the query set, where labels are propagated from support to query. Under transductive inference, the edge-labeling graph neural network (EGNN) proposed in [11] learns more accurate edge-labels through exploring the intra-cluster similarity and the inter-cluster dissimilarity. Our method takes the advantage of transductive inference that samples with higher prediction confidence can be obtained when the test samples are predicted as a whole.

Semi-Supervised Few-Shot Learning In semi-supervised few-shot learning, an extra unlabeled set not contained in current episode is used to improve classification accuracy [29, 25, 14]. In [25], the extended versions of Prototypical Networks [30] are proposed to use unlabeled data to create class prototypes by Soft k -Means. LST [14] employs pseudo-labeling strategy to the unlabeled set, then it re-trains and fine-tunes the base model based on the pseudo-labeled data. For recognizing the novel classes, it utilizes dynamically sampled data which is not contained in the current episode. Different from these methods, the unlabeled data in our method comes from the query set and we requires no extra datasets besides the support and query set.

3 Methodology

We firstly use cosine similarity based prototypical network (CSPN) to learn a discriminative feature space and get the basic prototypes of few-shot classes. Then we figure out two influencing factors in prototype computation: the intra-class bias and the cross-class bias. Accordingly, we propose the bias diminishing (BD) method for prototype rectification in transductive setting.

3.1 Denotation

At the training stage, a labeled dataset \mathcal{D} of base classes \mathcal{C}_{base} is given to train the feature extractor $F_{\theta}(\cdot)$ and the cosine classifier $C(\cdot|W)$. At the inference stage, we aim to recognize few-shot classes \mathcal{C}_{few} with K labeled images per class. Episodic sampling is adopted to form such N -way K -shot tasks. Each

episode consists of a support set \mathcal{S} and a query set \mathcal{Q} . In the support set, all samples x are labeled and we use the extracted features $X = F_\theta(x)$ to compute the prototypes P of few-shot classes. The samples in the query set are unlabeled for test.

3.2 Cosine Similarity Based Prototypical Network

We propose a metric learning based method: cosine similarity based prototypical network (CSPN) to compute the basic prototypes of few-shot classes. Training a good feature extractor that can extract discriminative features is of great importance. Thus, we firstly train a feature extractor $F_\theta(\cdot)$ with a cosine similarity based classifier $C(\cdot|W)$ on the base classes. The cosine classifier $C(\cdot|W)$ is:

$$C(F_\theta(x) | W) = \text{Softmax}(\tau \cdot \text{Cos}(F_\theta(x), W)) \quad (1)$$

where W is the learnable weight of the base classes and τ is a scalar parameter. We target to minimize the negative log-likelihood loss on the supervised classification task:

$$L(\theta, W | \mathcal{D}) = \mathbb{E}[-\log C(F_\theta(x) | W)] \quad (2)$$

At the inference stage, retraining $F_\theta(\cdot)$ and classification weights on the scarce data of \mathcal{C}_{few} classes is likely to run into overfitting. To avoid it, we directly compute the basic prototype P_n of class n as follows:

$$P_n = \frac{1}{K} \sum_{i=1}^K \bar{X}_{i,n} \quad (3)$$

where \bar{X} is the normalized feature of support samples. The query samples can be classified by finding the nearest prototype based on cosine similarity.

3.3 Bias Diminishing for Prototype Rectification

In CSPN, we can obtain the basic prototypes by simply averaging the features of support samples. However, the prototypes computed in such low-data regimes are biased against the expected prototypes we want to find. Therefore, we identify two influencing factors: the intra-class bias and the cross-class bias, and accordingly propose the bias diminishing approach.

The intra-class bias within a class is defined by Eq. (4):

$$B_{intra} = \mathbb{E}_{X' \sim p_{X'}}[X'] - \mathbb{E}_{X \sim p_X}[X] \quad (4)$$

where $p_{X'}$ is the distribution of all data belonging to a certain class and p_X is the distribution of the available labeled data of this class. It is easy to observe the difference between the expectations of the two distributions. The difference becomes more significant in low-data regimes. Since the prototype is computed by feature averaging, the intra-class bias also can be understood as the difference between the expected prototype and the actually computed prototype. The

expected prototype is supposed to be represented by the mean feature of all samples within a class. In practice, only a part of samples are available for training which is to say that, it is almost impossible to get the expected prototype. In few-shot scenario, we merely have K samples per few-shot class. The number of available samples are far less than the expected amount. Computed from scarce samples, the prototypes obviously tend to be biased.

To reduce the bias, we adopt the pseudo-labeling strategy to augment the support set, which assigns temporary labels to the unlabeled data according to their prediction confidence [14]. Pseudo-labeled samples can be augmented into the support set such that we can compute new prototypes in a ‘higher-data’ regime. We can simply select top Z confidently predicted query samples per class to augment the support set \mathcal{S} with their pseudo labels. We use CSPN as recognition model to get the prediction scores. Then we have an augmented support set with confidently predicted query samples: $\mathcal{S}' = \mathcal{S} \cup \mathcal{Q}_{pseudo}^Z$. Since some pseudo-labeled samples are likely to be misclassified, simple averaging with the same weights is possible to result in larger bias in prototype computation. To compute the new prototypes in a more reasonable way, we use the weighted sum of X' as the rectified prototype. We note that X' refers to the feature of the sample in \mathcal{S}' including both original support samples and pseudo-labeled query samples. The rectified prototype of a class is thus computed from the normalized features \bar{X}' :

$$P'_n = \sum_{i=1}^{Z+K} w_{i,n} \cdot \bar{X}'_{i,n} \quad (5)$$

where $w_{i,n}$ is the weight indicating the relation of the augmented support samples and the basic prototypes. The weight is computed by:

$$w_{i,n} = \frac{\exp(\varepsilon \cdot \text{Cos}(X'_{i,n}, P_n))}{\sum_{j=1}^{K+Z} \exp(\varepsilon \cdot \text{Cos}(X'_{j,n}, P_n))} \quad (6)$$

ε is a scalar parameter and P_n is the basic prototype obtained in Section 3.2. Samples with larger cosine similarity to the basic prototypes hold larger proportions in prototype rectification. Compared with the basic prototype P_n , the rectified prototype P'_n distributes closer to the expected prototype.

The cross-class bias refers to the distance between the mean vectors of support and query datasets. It is derived from the domain adaptation problem where the mean value is used as a type of the first order statistic information to represent a dataset [34]. Minimizing the distance between different domains is a typical method of mitigating domain gaps. Since the support set and the query set are assumed to distribute in the same domain, the distance between them is the distribution bias rather than the domain gap. The cross-class bias B_{cross} is formulated as:

$$B_{cross} = \mathbb{E}_{X_s \sim p_S} [X_s] - \mathbb{E}_{X_q \sim p_Q} [X_q] \quad (7)$$

where p_S and p_Q respectively represent the distributions of support and query sets. Notably, the support set \mathcal{S} and the query set \mathcal{Q} include N few-shot classes in Eq. (7). To diminish B_{cross} , we can shift the query set towards the support

set. In practice, we add a shifting term ξ to each normalized query feature \bar{X}_q and ξ is defined as:

$$\xi = \frac{1}{|\mathcal{S}|} \sum_{i=1}^{|\mathcal{S}|} \bar{X}_{i,s} - \frac{1}{|\mathcal{Q}|} \sum_{j=1}^{|\mathcal{Q}|} \bar{X}_{j,q} \quad (8)$$

The detailed derivation of ξ is given in the next section.

4 Theoretical Analysis

We give the theoretical analysis to show the rationality of our proposed bias diminishing method.

4.1 Lower Bound of the Expected Performance

We derive the formulation of our expected performance in theory and point out what factors influence the final result. We use X to represent the feature of a class. For clear illustration, the formulation of class prototype we use in this section is given:

$$P = \frac{\sum_i^T \bar{X}_i'}{T} \quad (9)$$

where $T = K + Z$, $\bar{X}_i' \in \mathcal{S}'$ and \mathcal{S}' is a subset sampled from X . \bar{X} is the normalized feature and \bar{P} is the normalized prototype. For cosine similarity based prototypical network, an expected prototype should have the largest cosine similarity to all samples within its class. Our objective is to maximize the expected cosine similarity which is positively correlated with the classification accuracy. It is formulated as:

$$\max \mathbb{E}_P[\mathbb{E}_X[\text{Cos}(P, X)]] \quad (10)$$

And we derive it as:

$$\begin{aligned} \mathbb{E}_P[\mathbb{E}_X[\text{Cos}(P, X)]] &= \mathbb{E}_{P,X}[\bar{P} \cdot \bar{X}] \\ &= \mathbb{E}[\bar{X}] \cdot \mathbb{E}\left[\frac{P}{\|P\|_2}\right] \end{aligned} \quad (11)$$

From previous works [21, 26], we know that:

$$\mathbb{E}\left[\frac{A}{B}\right] = \frac{\mathbb{E}[A]}{\mathbb{E}[B]} + O(n^{-1}) \quad (\text{first order}) \quad (12)$$

where A and B are random variables. In Eq. (12), $\frac{\mathbb{E}[A]}{\mathbb{E}[B]}$ is the first order estimator of $\mathbb{E}\left[\frac{A}{B}\right]$. Thus, Eq. (11) is approximate to:

$$\mathbb{E}_P[\mathbb{E}_X[\text{Cos}(P, X)]] \approx \frac{\mathbb{E}[\bar{X}] \cdot \mathbb{E}[P]}{\mathbb{E}[\|P\|_2]} \quad (13)$$

Based on Cauchy-Schwarz inequality, we have:

$$\mathbb{E}[\|P\|_2] \leq \sqrt{\mathbb{E}[\|P\|_2^2]} \quad (14)$$

P and \bar{X} are D -dimensional vectors which can be denoted as $P = [p_1, p_2, \dots, p_D]$ and $\bar{X} = [\bar{x}_1, \bar{x}_2, \dots, \bar{x}_D]$ respectively. In our method, we assume that each dimension of a vector is independent from each other. Then, we can derive that:

$$\begin{aligned} \mathbb{E}[\|P\|_2^2] &= \mathbb{E}[\sum_{i=1}^D p_i^2] = \sum_{i=1}^D [Var[p_i] + \mathbb{E}[p_i]^2] \\ &= \sum_{i=1}^D [Var[p_i] + \mathbb{E}[\bar{x}_i]^2] \\ &= \sum_{i=1}^D [\frac{1}{T} Var[\bar{x}_i] + \mathbb{E}[\bar{x}_i]^2] \end{aligned} \quad (15)$$

Thus, the lower bound of the expected cosine similarity is formulated as:

$$\begin{aligned} \mathbb{E}_P[\mathbb{E}_X[Cos(P, X)]] &\geq \frac{\mathbb{E}[\bar{X}] \cdot \mathbb{E}[P]}{\sqrt{\mathbb{E}[\|P\|_2^2]}} \\ &= \frac{\sum_{i=1}^D \mathbb{E}[\bar{x}_i]^2}{\sqrt{\frac{1}{T} \sum_{i=1}^D Var[\bar{x}_i] + \sum_{i=1}^D \mathbb{E}[\bar{x}_i]^2}} \end{aligned} \quad (16)$$

Maximizing the expected accuracy is approximate to maximize its lower bound of the cosine similarity as shown in Eq. (16). It can be seen that the number T of the sample is positively correlated with the lower bound of the expected performance. Thus, we import more pseudo-labeled samples into prototype computation. *The rationality of the pseudo-labeling strategy in improving few-shot accuracy is that, it can effectively raise the lower bound of the expected performance.*

4.2 Derivation of Shifting Term ξ

We propose to reduce the cross-class bias by feature shifting and the derivation of shifting term ξ is provided as follows. In N -way K -shot Q -query tasks, the accuracy can be formalized as:

$$Acc = \frac{1}{NQ} \sum_i^N \sum_q^Q \mathbb{1}(y_{i,q} == i) \quad (17)$$

$$= \frac{1}{NQ} \sum_i^N \sum_q^Q \mathbb{1}(Cos(P_i, X_{i,q}) > \max_{j \neq i} \{Cos(P_j, X_{i,q})\}) \quad (18)$$

where $y_{i,q}$ is the predicted label and i is the true class label. $\mathbb{1}(b)$ is an indicator function. $\mathbb{1}(b) = 1$ if b is true and 0 otherwise. P_i is the prototype of class i

and $X_{i,q}$ is the q -th query feature of class i . Based on Eq. (18), the accuracy formulation can be further rewritten as:

$$Acc = \frac{1}{NQ} \sum_i^N \sum_q^Q \mathbb{1}(Cos(P_i, X_{i,q}) > t_i) \quad (19)$$

where t_i denotes the cosine similarity threshold of the i -th class. Improving the accuracy is equal to maximize the cosine similarity $Cos(\cdot)$.

As mentioned above, there is a bias between the support and query set of a class i . We assume that the bias can be diminished by adding a shifting term ξ_i to the query samples. Since the class labels are unknown, we approximately add the same term ξ to all query samples. The term ξ should follow the objective:

$$\arg \max_{\xi} \frac{1}{NQ} \sum_i^N \sum_q^Q Cos(P_i, X_{i,q} + \xi) \quad (20)$$

We assume that each feature X can be represented as $X = P + \epsilon$. Eq. (20) can be further formalized as:

$$\arg \max_{\xi} \frac{1}{NQ} \sum_i^N \sum_q^Q Cos(P_i, P_i + \epsilon_{i,q} + \xi) \quad (21)$$

To maximize the cosine similarity, we should minimize the following objective:

$$\min \frac{1}{NQ} \sum_i^N \sum_q^Q (\epsilon_{i,q} + \xi) \quad (22)$$

The term ξ is thus computed:

$$\xi = -E[\epsilon] \quad (23)$$

$$= \frac{1}{NQ} \sum_i^N \sum_q^Q (P_i - X_{i,q}) \quad (24)$$

We can see that Eq. (24) is in line with Eq. (8). For cosine similarity computation, the shifting term is calculated from the normalized features as displayed in Section 3.3.

5 Experiments

5.1 Datasets

miniImageNet consists of 100 randomly chosen classes from ILSVRC-2012 [27]. We adopt the split proposed in [24] where the 100 classes are split into 64 training classes, 16 validation classes and 20 test classes. Each class contains 600 images of size 84×84 . **tieredImageNet** [25] is also a derivative of ILSVRC-2012 [27] containing 608 low-level categories, which are split into 351, 97, 160 categories for training, validation, test with image size of 84×84 . **Meta-Dataset** [32] is a new benchmark that is large-scale and consists of diverse datasets for training and evaluating models.

5.2 Implementation details

We train the base recognition model CSPN in the supervised way with SGD optimizer and test the validation set on 5-way 5-shot tasks for model selection. WRN-28-10 [35] is used as the main backbone. ConvNets [8] and ResNet-12 [13] are used for ablation. The results are averaged from 600 randomly sampled episodes. Each episode contains 15 query samples per class. The initial value of τ is 10 and ε is fixed at 10. More details are shown in the supplementary materials.

Table 1. Average accuracy (%) comparison on miniImageNet. † Training set and validation set are used for training.

Setting	Methods	Backbone	miniImageNet	
			1-shot	5-shot
Inductive	Matching Network [33]	ConvNet-64	43.56±0.84	55.31±0.73
	MAML [6]	ConvNet-32	48.70±1.84	63.11±0.92
	Prototypical Networks† [30]	ConvNet-64	49.42±0.78	68.20±0.66
	Relation Net [31]	ConvNet-256	50.44±0.82	65.32±0.70
	SNAIL [18]	ResNet-12	55.71±0.99	68.88±0.92
	LwoF [8]	ConvNet-128	56.20±0.86	73.00±0.64
	AdaResNet [19]	ResNet-12	56.88±0.62	71.94±0.57
	TADAM [22]	ResNet-12	58.50±0.30	76.70±0.30
	Activation to Parameter† [23]	WRN-28-10	59.60±0.41	73.74±0.19
	LEO† [28]	WRN-28-10	61.76±0.08	77.59±0.12
	MetaOptNet-SVM [13]	ResNet-12	62.64±0.61	78.63±0.46
	BFSL [7]	WRN-28-10	62.93±0.45	79.87±0.33
Semi-Supervised	ML [25]	ConvNet-128	49.04±0.31	62.96±0.14
	LST [14]	ResNet-12	70.1±1.9	78.7±0.8
Transductive	TPN [16]	ConvNet-64	55.51±0.86	69.86±0.65
	EGNN [11]	ConvNet-256	-	76.37
	Transductive Fine-Tuning [4]	WRN-28-10	65.73±0.68	78.40±0.52
	BD-CSPN (ours)	WRN-28-10	70.31±0.93	81.89±0.60

5.3 Results on miniImageNet and tieredImageNet

The results on miniImageNet and tieredImageNet are shown in Table 1 and Table 2 respectively. It can be seen that we achieve state-of-the-art performance in all cases. Compared with existing transductive methods [16, 11, 4], our proposed BD-CSPN consistently achieves the best performance on both datasets. EGNN [11] transductively learns edge-labels through exploring the intra-cluster similarity and the inter-cluster dissimilarity. Transductive Fine-Tuning [4] is newly published, providing a strong baseline by simple fine-tuning techniques. In comparison with TPN [16], we achieve better results with a simpler implementation of label propagation technique. Given the similar backbone ConvNet-128

Table 2. Average accuracy (%) comparison on tieredImageNet. * Results by our implementation. ‡ Training set and validation set are used for training.

Setting	Methods	Backbone	tieredImageNet	
			1-shot	5-shot
Inductive	MAML [6]	ConvNet-32	51.67±1.81	70.30±1.75
	Prototypical Networks‡ [30]	ConvNet-64	53.31±0.89	72.69±0.74
	Relation Net [31]	ConvNet-256	54.48±0.93	71.32±0.78
	LwoF [8]	ConvNet-128	60.35±0.88*	77.24±0.72*
	LEO‡ [28]	WRN-28-10	66.33±0.05	81.44±0.09
	MetaOptNet-SVM [13]	ResNet-12	65.99±0.72	81.56±0.53
Semi-Supervised	ML [25]	ConvNet-128	51.38±0.38	69.08±0.25
	LST [14]	ResNet-12	77.7±1.6	85.2±0.8
Transductive	TPN [16]	ConvNet-64	59.91±0.94	73.30±0.75
	EGNN [11]	ConvNet-256	-	80.15
	Transductive Fine-Tuning [4]	WRN-28-10	73.34±0.71	85.50±0.50
	BD-CSPN (ours)	WRN-28-10	78.74±0.95	86.92±0.63

on miniImageNet, BD-CSPN produces good results of 61.74% and 76.12% on 1-shot and 5-shot tasks respectively, surpassing TPN by large margins.

Our method also shows superiority compared with existing semi-supervised methods [25, 14]. Note that LST [14] uses extra unlabeled data as auxiliary information in evaluation, which is not contained in current episode. It re-trains and fine-tunes the model on each novel task. We have a simpler technique without re-training and fine-tuning which is more efficient in computation.

5.4 Results on Meta-Dataset

To further illustrate the effectiveness of our method, we show 5-shot results on the newly proposed Meta-Dataset [32] in Table 3. The average rank of our 5-shot model is **1.9**. More details are provided in our supplementary materials.

Table 3. 5-shot results on Meta-Dataset: the model is trained on ILSVRC-2012 only and test on the listed test sources.

	Test 5-shot	Test 5-shot	Test 5-shot	Test 5-shot	Test 5-shot				
ILSVRC	59.80	Omniglot	78.29	Aircraft	43.42	Birds	67.22	Textures	54.82
Quick Draw	58.80	Fungi	61.56	VGG Flower	83.88	Traffic Signs	68.68	MSCOCO	52.69

5.5 Ablation Study

The ablation results are shown in Table 4. We display the results of CSPN as baselines which are obtained in inductive setting. The network is trained on traditional supervised tasks (64-way), following the setting in [8, 7]. It achieves

better performance than some complicated meta-trained methods [23, 28] with the same backbone, as shown in Table 1. Based on CSPN, our BD module makes an improvement by large margins up to 9% and 3% on 1-shot and 5-shot tasks respectively. It leads to relatively minor improvements in 5-shot scenarios.

Table 4. Ablative results of bias diminishing module. CSPN: without bias diminishing modules; BD_c -CSPN: with cross-class bias diminishing module; BD_i -CSPN: with intra-class bias diminishing module; BD -CSPN: with both modules.

Dataset	Method	1-shot	5-shot	Dataset	Method	1-shot	5-shot
	CSPN	61.84	78.64		CSPN	69.20	84.31
miniImageNet	BD_c -CSPN	62.54	79.32	tieredImageNet	BD_c -CSPN	70.84	84.99
	BD_i -CSPN	69.81	81.58		BD_i -CSPN	78.12	86.67
	BD -CSPN	70.31	81.89		BD -CSPN	78.74	86.92

Ablation of Intra-Class Bias Diminishing It can be seen in Table 4 that BD_i -CSPN brings in significant improvements on both datasets. The intra-class bias diminishing module especially shows its merit in 1-shot scenarios. With intra-class bias diminished, the accuracy on 1-shot miniImageNet increases from 61.84% to 69.81% and the accuracy on 1-shot tieredImageNet raises to 78.12% from 69.20%.

Furthermore, to intuitively demonstrate the influence of our proposed intra-class bias diminishing module, we display the 5-way accuracy in Fig. 2(a)-2(b). The results are reported without using cross-class bias diminishing module. It shows a coincident tendency that with more pseudo-labeled samples, there is an obvious growth of classification accuracy. We use the validation set to determine the value of Z and set it to 8 for accuracy comparison in Table 1 and Table 2.

Theoretical Value As we know, the expected accuracy $Acc(P, X)$ has a positive correlation with the expected cosine similarity. Then we derive the first-order estimation of $Acc(P, X)$ from Eq. (16) which is formulated as:

$$Acc(P, X) \approx \eta \cdot \frac{\alpha}{\sqrt{\lambda \cdot \frac{1}{K+Z} + \alpha}} \quad (25)$$

where η is a coefficient and $K + Z = T$. λ and α are values correlated with the variance term and the expectation term in Eq. (16). The theoretical values of λ and α can be approximately computed from the extracted features. Furthermore, we can compute the value of η by 1-shot and 5-shot accuracies of CSPN. Thus, the number Z is the only variable in Eq. (25). The theoretical curves are displayed as the dashed lines in Fig. 2(c) to show the impact of Z on classification accuracy. The dashed lines, showing the theoretical lower bound of the expected accuracy, have a consistent tendency with our experiment results in Fig. 2(a)-2(b). Since the cosine similarity is continuous and the accuracy is discrete, the accuracy stops increasing when the cosine similarity grows to a certain value.

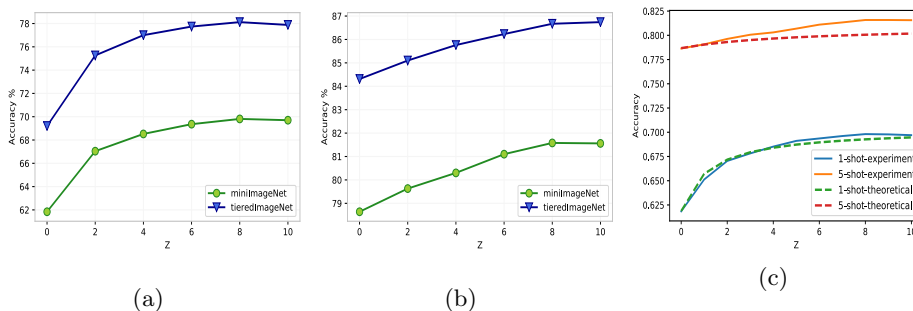


Fig. 2. Effectiveness of intra-bias diminishing. Z : the number of pseudo-labeled samples. (a) 5-way 1-shot results. (b) 5-way 5-shot results. (c) Theoretical value on miniImageNet. The experiment results (solid lines) show a consistent tendency with the theoretical results (dashed lines).

T-SNE Visualization We show t-SNE visualization of our intra-bias diminishing method in Fig. 3(a) for intuitive illustration. The basic prototype of each class is computed from the support set while the rectified prototype is computed from the augmented support set. In this section, the expected prototype refers to the first term in Eq. (4) which is represented by the average vector of all samples (both support and query samples) of a class in an episode. Due to the scarcity of labeled samples, there is a large bias between the basic prototype and the expected prototype. The bias can be reflected by the distance between the stars and the triangles in Fig. 3(a).

Ablation of Cross-Class Bias Diminishing Table 4 shows the ablative results of the cross-class bias diminishing module. It illustrates an overall improvement as a result of diminishing the cross-class bias. Moving the whole query set towards the support set center by importing the shifting term ξ is an effective approach to reduce the bias between the two datasets. For example, the accuracy increases by 1.64% on 1-shot tieredImageNet.

T-SNE Visualization In few-shot learning, the support set includes far less samples compared with the query set in an episode. There exists a large distance between the two mean vectors of the datasets. We aim to decrease the distance by shifting the query samples towards the center of the support set as shown in Fig. 3(b). It depicts the spatial changing of the query samples, before and after cross-class bias diminishing. The significant part is zoomed in for clear visualization, where the query samples with BD_{cross} (marked in green) distribute more closely to the center of support set.

Ablation of Backbone The results on miniImageNet are displayed in Table 5 and more ablation results are given in the supplementary materials. Our method also shows good performance based on ConvNet-128 and ResNet-12, which is

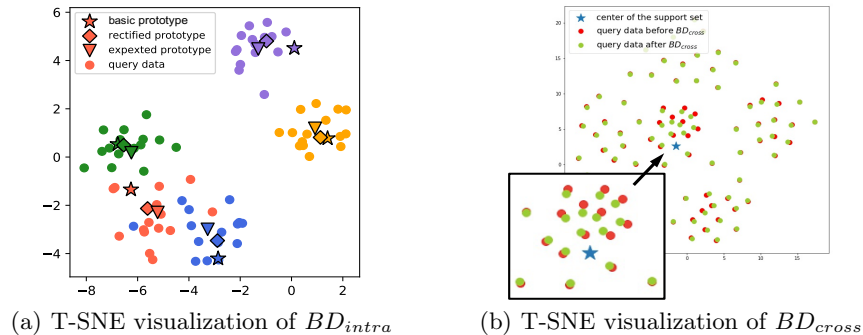


Fig. 3. We randomly sample a 5-way 1-shot episode on tieredImageNet. Different classes are marked in different colors. Best viewed in color with zoom in.

Table 5. Ablation of **backbones** and result **comparison with TFT** (Transductive Fine-Tuning [4]) on miniImageNet. * The backbone is ConvNet-64.

1-shot	CSPN	BD-CSPN	TFT	5-shot	CSPN	BD-CSPN	TFT
ConvNet-128	55.62	61.74	50.46*	ConvNet-128	72.57	76.12	66.68*
ResNet-12	59.14	65.94	62.35	ResNet-12	76.26	79.23	74.53
WRN-28-10	61.84	70.31	65.73	WRN-28-10	78.64	81.89	78.40

better than most approaches in Table 1. For example, with ResNet-12, we achieve 79.23% in 5-shot scenario, outperforming the strongest baselines: 78.7% [14] and 78.63% [13].

5.6 Comparison with Transductive Fine-Tuning

We compare our method with TFT [4] in Table 5, which is recently proposed as a new baseline for few-shot image classification. BD-CSPN outperforms it given different backbones. For example, we achieve better results which are higher than TFT by 3% to 5% given ResNet-12. Since BD-CSPN and TFT conduct experiments in the same transductive setting, the comparison between these two methods is more persuasive to demonstrate the effectiveness of the approach.

6 Conclusions

In this paper, we propose a powerful method of prototype rectification in few-shot learning, which is to diminish the intra-class bias and the cross-class bias of class prototypes. Our theoretical analysis verifies that, the proposed bias diminishing method is effective in raising the lower bound of the expected performance. Extensive experiments on three few-shot benchmarks demonstrate the effectiveness of our method. The proposed bias diminishing method achieves significant improvements in transductive setting by large margins (e.g. 8.47% on 1-shot miniImageNet and 9.54% on 1-shot tieredImageNet).

References

1. Allen, K., Shelhamer, E., Shin, H., Tenenbaum, J.: Infinite mixture prototypes for few-shot learning. In: ICML. pp. 232–241 (2019)
2. Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M.W., Pfau, D., Schaul, T., Shillingford, B., de Freitas, N.: Learning to learn by gradient descent by gradient descent. In: NIPS. pp. 3981–3989 (2016)
3. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C.F., Huang, J.B.: A closer look at few-shot classification. In: ICLR (2019)
4. Dhillon, G.S., Chaudhari, P., Ravichandran, A., Soatto, S.: A baseline for few-shot image classification. In: ICLR (2020)
5. Fei-Fei, L., Fergus, R., Perona, P.: One-shot learning of object categories. vol. 28, pp. 594–611 (2006)
6. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML. pp. 1126–1135 (2017)
7. Gidaris, S., Bursuc, A., Komodakis, N., Perez, P.P., Cord, M.: Boosting few-shot visual learning with self-supervision. In: ICCV. pp. 8058–8067 (2019)
8. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: CVPR. pp. 4367–4375 (2018)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
10. Joachims, T.: Transductive inference for text classification using support vector machines. In: ICML. pp. 200–209 (1999)
11. Kim, J., Kim, T., Kim, S., Yoo, C.D.: Edge-labeling graph neural network for few-shot learning. In: CVPR. pp. 11–20 (2019)
12. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. vol. 141, pp. 1097–1105 (2012)
13. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: CVPR. pp. 10657–10665 (2019)
14. Li, X., Sun, Q., Liu, Y., Zhou, Q., Zheng, S., Chua, T.S., Schiele, B.: Learning to self-train for semi-supervised few-shot classification. In: NeurIPS (2019)
15. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few shot learning. (2017)
16. Liu, Y., Lee, J., Park, M., Kim, S., Yang, E., Hwang, S.J., Yang, Y.: Learning to propagate labels: Transductive propagation network for few-shot learning. In: ICLR (2019)
17. Miller, E., Matsakis, N., Viola, P.: Learning from one example through shared densities on transforms. In: CVPR. vol. 1, pp. 464–471 (2000)
18. Mishra, N., Rohaninejad, M., Chen, X., Abbeel, P.: A simple neural attentive meta-learner. In: ICLR (2018)
19. Munkhdalai, T., Yuan, X., Mehri, S., Trischler, A.: Rapid adaptation with conditionally shifted neurons. In: ICML. pp. 3661–3670 (2018)
20. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. (2018)
21. Nowozin, S.: Optimal decisions from probabilistic models: The intersection-over-union case. In: CVPR. pp. 548–555 (2014)
22. Oreshkin, B.N., Lopez, P.R., Lacoste, A.: Tadam: Task dependent adaptive metric for improved few-shot learning. In: NIPS. pp. 721–731 (2018)
23. Qiao, S., Liu, C., Shen, W., Yuille, A.L.: Few-shot image recognition by predicting parameters from activations. In: CVPR. pp. 7229–7238 (2018)

24. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2017)
25. Ren, M., Ravi, S., Triantafillou, E., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: ICLR (2018)
26. Rice, S.H.: The expected value of the ratio of correlated random variables. Texas Tech University (2015)
27. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. vol. 115, pp. 211–252 (2015)
28. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: ICLR (2019)
29. Satorras, V.G., Estrach, J.B.: Few-shot learning with graph neural networks. In: ICLR (2018)
30. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. In: NIPS. pp. 4077–4087 (2017)
31. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR. pp. 1199–1208 (2018)
32. Triantafillou, E., Zhu, T., Dumoulin, V., Lamblin, P., Evci, U., Xu, K., Goroshin, R., Gelada, C., Swersky, K., Manzagol, P.A., Larochelle, H.: Meta-dataset: A dataset of datasets for learning to learn from few examples. In: ICLR (2020)
33. Vinyals, O., Blundell, C., Lillicrap, T.P., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: NIPS. pp. 3637–3645 (2016)
34. Wang, Y., Li, W., Dai, D., Gool, L.V.: Deep domain adaptation by geodesic distance minimization. In: ICCVW. pp. 2651–2657 (2017)
35. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: BMVC (2016)
36. Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Schlkopf, B.: Learning with local and global consistency. In: NIPS. pp. 321–328 (2003)