

Cooperative Bi-path Metric for Few-shot Learning

Zeyuan Wang¹, Yifan Zhao¹, Jia Li^{1,3,4*}, Yonghong Tian^{2,4}

¹State Key Laboratory of Virtual Reality Technology and Systems, SCSE, Beihang University, Beijing, 100191, China

²School of Electronics Engineering and Computer Science, Peking University, Beijing, 100871, China

³Beijing Advanced Innovation Center for Big Data and Brain Computing, Beihang University, Beijing, 100191, China

⁴Peng Cheng Laboratory, Shenzhen, 518066, China

ABSTRACT

Given base classes with sufficient labeled samples, the target of few-shot classification is to recognize unlabeled samples of novel classes with only a few labeled samples. Most existing methods only pay attention to the relationship between labeled and unlabeled samples of novel classes, which do not make full use of information within base classes. In this paper, we make two contributions to investigate the few-shot classification problem. First, we report a simple and effective baseline trained on base classes in the way of traditional supervised learning, which can achieve comparable results to the state of the art. Second, based on the baseline, we propose a cooperative bi-path metric for classification, which leverages the correlations between base classes and novel classes to further improve the accuracy. Experiments on two widely used benchmarks show that our method is a simple and effective framework, and a new state of the art is established in the few-shot classification field.

CCS CONCEPTS

• **Computing methodologies** → **Supervised learning by classification**; Dimensionality reduction and manifold learning; • **Information systems** → *Similarity measures*.

KEYWORDS

few-shot learning, image classification, metric learning, locally linear embedding

ACM Reference Format:

Zeyuan Wang, Yifan Zhao, Jia Li, Yonghong Tian. 2020. Cooperative Bi-path Metric for Few-shot Learning. In *Proceedings of the 28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394171.3413946>

1 INTRODUCTION

Image recognition provides an intuitive and fundamental way to understand the visual world, which is also deeply rooted in many advanced kinds of research, including autonomic production and

*Jia Li is the corresponding author (E-mail: jjiali@buaa.edu.cn). Website: <http://cvteam.net>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '20, October 12–16, 2020, Seattle, WA, USA
 © 2020 Association for Computing Machinery.
 ACM ISBN 978-1-4503-7988-5/20/10...\$15.00
<https://doi.org/10.1145/3394171.3413946>

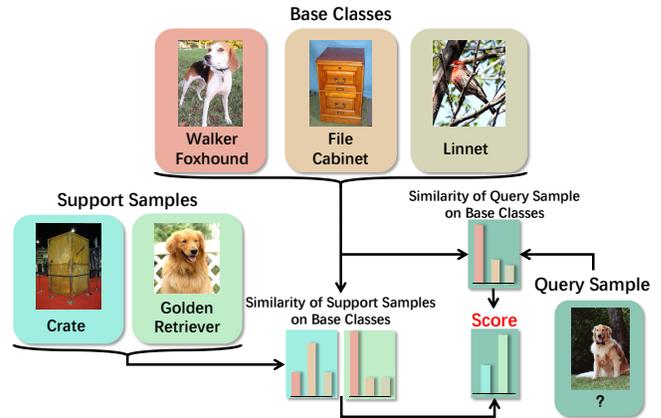


Figure 1: Motivation of the proposed approach.

artificial intelligence. As the crucial part of visual content, image recognition has witnessed a great progress with the proposals of large dataset [1] and deep learning techniques [2, 3]. However, the notable performance of these methods heavily relies on the large amount of manually annotated datasets, which are labour-intensive and sometimes inaccessible, e.g., ImageNet [1] with over 15 million annotations. Remarkably, humans and other animals seem to have the potential to recognize one identity with very less related knowledge. Therefore, the few-shot learning (FSL) problem with limited seen knowledge forces the model to make a typical generalization of each class, which is a more realistic setting in some extreme industrial applications.

By training a model on base classes that contain sufficient labeled samples, the goal of few-shot learning is to make the model generalize well on the novel classes which do not intersect with the base classes, i.e., correctly classifying unlabeled samples (*query samples*) according to a small number of labeled samples (*support samples*). To make the conditions of the training phase match those of the testing phase, Matching Networks [4] firstly suggested that both training and testing should adopt *episodic* procedure, which came from meta-learning. Models will meet many few-shot learning tasks in both the training and testing phases. Each task consists of several classes, and each class contains a few support samples and several query samples. Leading by the pioneer work [4], many subsequent researches [5–17] followed this episodic learning and achieved notable improvements. However, some recent works [18–21] did not follow this factitious sampling setting, but directly trained model in the way of traditional supervised learning. Thus a natural concern arises, **is episodic training necessary for few-shot learning?** Keeping this in our mind, we first make extensive

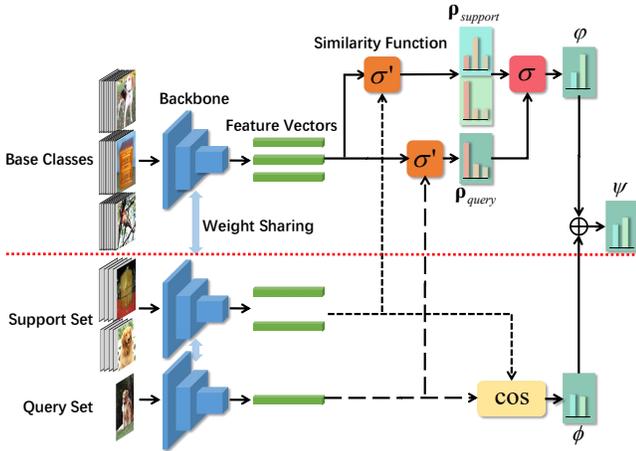


Figure 2: Illustration of the pipeline of Cooperative Bi-path Metric. The top half part upon the red dotted line represents our proposed transductive similarity, and the bottom half part is the classic inductive similarity. $\rho_{support}$ refers to the similarity distributions of support set on base classes, and ρ_{query} refers to those of query set. The final classification score ψ is the weighted sum of ϕ and ϕ , where ϕ is the conventional cosine similarity between support set and query set, and ϕ is the similarity between $\rho_{support}$ and ρ_{query} .

experimental analyses on two commonly used benchmarks [4, 12]. Counter-intuitively, without the always-used episodic training process, state-of-the-art performance can also be achieved through adequate training strategy with all the samples in base classes. **This finding not only brings us a rethinking of this conventional setting but also can be considered as a high-performance baseline for few-shot learning.**

To start from another perspective, metric learning [22–33] is a major genre in the field of few-shot learning. This kind of method classifies query samples by learning a feature extractor on the base classes, extracting features of samples of the novel classes during testing, and measuring the distance or similarity between labeled support samples and unlabeled query samples. However, most of the existing metric learning methods [4, 9, 11, 14–17, 20] for few-shot learning focused on the correlation between support samples and query samples within novel classes, and did not take full use of the information of base classes. Unlike the above-mentioned approaches, we make use of the base classes to assist in classifying query samples of novel classes during the testing phase. **Our motivation is that for samples with identical labels of novel classes, their similarity distributions on base classes are more likely to be similar, while different classes’ samples should have different distributions from each other.** For example, as shown in the Fig. 1, the *golden retrievers* are more similar to *walker foxhounds* than *file cabinets*, meanwhile the *crates* are more similar to *file cabinets* than *walker foxhounds*. We emphasize this kind of relationship information between the base classes and the novel classes is valuable and helpful for classifying unlabeled samples. Inspired by this finding, we propose a novel metric named

Cooperative Bi-path Metric, as shown in Fig. 2, which not only considers the inductive similarity between support set and query set in the novel classes but also measures their transductive similarity distributions through base classes and further takes **the similarity of their similarity distributions on base classes** into consideration.

Our main contribution is three-fold: 1) We make extensive experimental analyses of the conventional episodic training and full-supervision training for the few-shot classification problem, and propose a new high-performance baseline based on this experimental finding. 2) We propose a novel Cooperative Bi-path Metric learning approach, which first exploits the base classes as an intermediary for facilitating the classification process. 3) We make extensive experimental analyses to demonstrate our findings. Compared with the existing methods, our proposed approach achieves new state-of-the-art results on two widely used benchmarks, *i.e.*, miniImageNet [4] and tieredImageNet [12].

2 RELATED WORK

There are two main branches of methods in the few-shot learning field, one is meta-learning based methods and the other is metric-learning based methods. The former adopts the episodic training procedure in the training stage and expects to learn the common attributes between different tasks through this procedure, namely the appropriate hyperparameters for models. The latter focuses on how to better extract features from samples and classify samples according to the extracted features, that is, the network is expected to learn a function that can properly measure the similarity between samples.

Meta-learning based methods. The purpose of meta-learning, or learning to learn [34–36], is to train a *meta-learner* to learn task-agnostic knowledge (or hyper-parameters), which can assist the training of learners on different tasks. Meta-learning based methods [5–7, 10, 18, 19] are a major branch in the field of few-shot learning. For example, Meta-LSTM [6] trained an LSTM-based [37] meta-learner to discover good initialization for learner’s parameters, as well as a mechanism for updating the learner’s parameters by a small sample set. Similarly, Meta-SGD [5] trained a meta-learner to produce learner’s initialization, update direction, and learning rate, but in a single meta-learning process. Also, MAML [7] aimed to find the appropriate initial parameters of the learner, so that the learner could converge rapidly with a few samples. Similar to MAML [7], Reptile [10] was a meta-learning method for finding neural network initialization parameters but simply performed SGD on each task without computing gradient twice as MAML [7] did. Gidaris *et al.* [18] and Qiao *et al.* [19] both used base classes to train a feature extractor in the first stage, then used the base classes training a classification weight generator (meta-learner) in the second stage, and used the classification weight generator and support samples to generate classification weight vectors (parameters of the learner) for novel classes during the testing phase. As we can see, all the above meta-learning methods require training a meta-learner and more or less fine-tuning on novel classes, which undoubtedly increases the complexity of the methods. However, our baseline is trained in the way of supervised learning, which can achieve comparable

performance to the state of the art, meanwhile greatly simplifies the training process.

Metric-learning based methods. Another branch in the few-shot learning field is metric-learning based methods [4, 8, 9, 11, 12, 14], which focus on embedding samples into a metric space so that the samples can be classified according to similarity or distance between each other. Matching Networks [4] used LSTM [37] and variants to extract feature vectors from support samples and query samples, and then classified query samples by calculating cosine similarities between support and query samples. Prototypical Networks [9] took the average of the feature vectors of support samples within a class as the prototype of the class, and assigned the query samples to the nearest prototype in Euclidean distance, meanwhile its convolutional neural network was trained end-to-end. Triantafyllou *et al.* [8] proposed an information retrieval-inspired approach viewing each batch point as a query that ranked the remaining ones and defined a model to optimize mean Average Precision over these rankings. Relation Network [11] adapted the same convolutional neural network as Prototypical Networks [9] to extract the features of support samples and query samples, but features of support samples and query samples were concatenated and input into nonlinear relation module to obtain classification scores. Ren *et al.* [12] presented a new problem: semi-supervised few-shot classification, in which support samples consisted of labeled and unlabeled samples, and proposed several novel extensions on Prototypical Networks [9]. Also, TADAM [14] learned a task-dependent metric with Metric Scaling, Task Conditioning, and Auxiliary Task Co-training for few-shot classification. Although the above metric-learning based methods could achieve admirable results on few-shot classification problem, they only considered the direct relationship between query samples and support samples when predicting the labels of query samples of novel classes, and ignored the relationship between novel classes and base classes. After observing this, we propose a novel Cooperative Bi-path Metric to utilize the relationship information between the support samples, query samples, and base classes to assist in the classification of queries, which can further improve the accuracy of classification.

3 METHOD

3.1 Problem Definition

Given a training set \mathbb{D}_{base} containing samples of base classes \mathbb{C}_{base} , the goal of few-shot learning is to train a model with \mathbb{D}_{base} to achieve high accuracy on the classification tasks obtained by sampling on test set \mathbb{D}_{novel} , which contains samples of novel classes \mathbb{C}_{novel} . The base classes are totally different from novel classes, that is, $\mathbb{C}_{base} \cap \mathbb{C}_{novel} = \emptyset$. And each task is composed of a support set $\mathbb{D}_{support}$ with labeled samples and a query set \mathbb{D}_{query} with unlabeled samples. For a N -way K -shot task, $\mathbb{D}_{support} = \left\{ \left(\mathbf{x}^{(i)}, y^{(i)} \right) \right\}_{i=1}^{N \times K}$ contains N classes and K support samples for each class. A model is trained to predict the labels of the query samples in \mathbb{D}_{query} as accurately as possible according to $\mathbb{D}_{support}$.

3.2 A Strong Baseline for Few-shot Learning

Is episodic training necessary? In the view of episodic training, a lot of works [4–17] made the procedure of training and testing

consistent to achieve higher performance. During training phase, tasks were sampled, then loss function of the model was calculated based on the tasks, and the network parameters were updated through the back propagation. However, some works [18–21] did not follow this setting, but trained classification networks in the way of traditional supervised learning.

To make fair comparisons with the previous works [13, 14, 16, 17, 20], we adopt ResNet-12 [3] as the baseline’s backbone and use the training set \mathbb{D}_{base} to optimize it in a fully-supervised manner. However, it is interesting to find that a network with only a supervised training strategy can also get superior classification results on novel classes (elaborated in Sect. 4). This indicates the basic knowledge learnt from base classes can not be further improved by the episodic learning and motivates the proposal of our strong baseline.

Baseline for few-shot learning. For the problem of few-shot learning, we advocate three meaningful cues in constructing a strong baseline: 1) Data augmentation: following prior work [17], we use horizontal flip, random crop and random erasing [38] as data augmentation. 2) Temperature in learning: inspired by [39], we also introduce a hyper-parameter called temperature, which was first applied in model distillation to change the smoothness of distribution after softmax normalization and the value of cross-entropy. 3) Dense classification: instead of embedding the image features as a vector, we apply dense classification loss [20] to regularize our model, *i.e.*, all the local feature vectors of the feature map before the last fully-connected layer are classified through the fully-connected layer without average pooling. For each training sample $(\mathbf{x}, y) \in \mathbb{D}_{base}$, the proposed baseline with loss \mathcal{L} has the following form:

$$\mathcal{L} = -\frac{1}{r} \sum_{i=1}^r \log \frac{\exp \left(t \left(\mathbf{f}^{(i)\top} \mathbf{p}^{(y)} + \mathbf{b}_y \right) \right)}{\sum_{j=1}^{|\mathbb{C}_{base}|} \exp \left(t \left(\mathbf{f}^{(i)\top} \mathbf{p}^{(j)} + \mathbf{b}_j \right) \right)}, \quad (1)$$

where t is the temperature hyperparameter. $|\cdot|$ is the cardinality of a set. $\mathbf{f}^{(i)} \in \mathbb{R}^c$ is the local vector at position i of the training sample’s feature map $\mathbf{F} \in \mathbb{R}^{c \times r}$ with channel dimension c and spatial resolution r . $\mathbf{p}^{(j)} \in \mathbb{R}^c$ is the parameter vector for class j in the fully-connected layer’s parameter matrix $\mathbf{P} \in \mathbb{R}^{c \times |\mathbb{C}_{base}|}$. \mathbf{b}_j is the bias for class j in the fully-connected layer’s bias vector $\mathbf{b} \in \mathbb{R}^{|\mathbb{C}_{base}|}$.

During testing phase, for a N -way K -shot task, a query sample in \mathbb{D}_{query} is assigned to the class \hat{y} with maximum classification score $\phi^{(n)}$:

$$\hat{y} = \arg \max_n \left(\phi^{(n)} \right). \quad (2)$$

Classification score $\phi^{(n)}$ for novel class n is defined as:

$$\phi^{(n)} = \cos \left(\mathbf{q}, \mathbf{s}^{(n)} \right). \quad (3)$$

And $\cos(\cdot, \cdot)$ is cosine similarity between two vectors. \mathbf{q} and $\mathbf{s}^{(n)}$ are feature vectors of query sample and class n respectively:

$$\begin{aligned} \cos(\mathbf{a}, \mathbf{b}) &= \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}, \\ \mathbf{q} &= \text{GAP}(\mathbf{Q}), \\ \mathbf{s}^{(n)} &= \frac{1}{K} \sum_{k=1}^K \text{GAP}(\mathbf{S}^{(n,k)}). \end{aligned} \quad (4)$$

And $\|\cdot\|$ is the L_2 norm of a vector. \mathbf{Q} and $\mathbf{S}^{(n,k)}$ are feature maps of the query sample and the k -th support sample of class n in $\mathbb{D}_{\text{support}}$. $\text{GAP}(\cdot)$ is the global average pooling on a feature map \mathbf{F} defined as:

$$\text{GAP}(\mathbf{F}) = \frac{1}{r} \sum_i^r \mathbf{f}^{(i)}, \quad (5)$$

where $\mathbf{f}^{(i)} \in \mathbb{R}^c$ is the local vector at position i of the feature map $\mathbf{F} \in \mathbb{R}^{c \times r}$.

3.3 Cooperative Bi-path Metric

As shown in Eq. (2) and (3), the previously proposed baseline in this paper as well as previous methods [4, 9, 11, 14–17, 20] simply classify query samples only according to support samples. The main drawback is that the prior knowledge on base classes is not fully exploited, which can also be complementary to classification decision. Thus a natural thought arises: the similarity distributions on base classes of support samples and query samples within the same class should also be similar, as illustrated in Fig. 1, and this information is useful for classifying query samples.

Starting from this point, we propose a novel method namely Cooperative Bi-path Metric as the classification criterion during the testing phase, which is shown in Fig. 2. Cooperative Bi-path Metric utilizes base classes as an intermediate way to assist with the classification of query samples. Our proposed metric measures the similarity by two individual paths: inductive similarity ϕ and transductive similarity φ . Most existing methods regard the former one as the only classification criterion, as shown in the lower half of Fig. 2, which calculates the inductive similarity ϕ (e.g. cosine similarity) between the support set and the query set. While Cooperative Bi-path Metric not only measures the inductive similarity ϕ but also uses base classes as an agent to calculate the transductive similarity φ between the support set and the query set, as shown in the upper part of Fig. 2. Firstly, it calculates the similarity distribution ρ_{support} and ρ_{query} of support set and query set on base classes, and then calculates the similarity between ρ_{support} and ρ_{query} , i.e., the transductive similarity φ between support set and query set. The final classification score ψ during the test phase is a weighted sum of ϕ and φ :

$$\begin{aligned} \hat{y} &= \arg \max_n \left(\psi^{(n)} \right), \\ \psi^{(n)} &= \alpha \phi^{(n)} + (1 - \alpha) \varphi^{(n)}, \\ \varphi^{(n)} &= \sigma \left(\rho_{\text{query}}^{(n)} \rho_{\text{support}}^{(n)} \right). \end{aligned} \quad (6)$$

And $\psi^{(n)}$ is Cooperative Bi-path Metric's final classification score for novel class n . $\phi^{(n)}$ is defined in Eq. (3). α is a hyperparameter to adjust the weight between $\phi^{(n)}$ and $\varphi^{(n)}$. $\sigma(\cdot)$ is a similarity

function that measures the similarity between two distributions, and it can be cosine similarity or negative Euclidean distance and so on. ρ_{query} and $\rho_{\text{support}}^{(n)}$ can be formally represented as:

$$\begin{aligned} \rho_{\text{query}} &= \sigma'(\mathbf{q}, \mathbf{B}), \\ \rho_{\text{support}}^{(n)} &= \sigma'(\mathbf{s}^{(n)}, \mathbf{B}), \end{aligned} \quad (7)$$

where \mathbf{q} and $\mathbf{s}^{(n)}$ are defined in Eq. (4). $\sigma'(\cdot, \cdot)$ is another similarity function that measures the similarity between a vector and each column of a matrix, while it can be similar to or different from $\sigma(\cdot, \cdot)$. \mathbf{B} is a feature matrix of base classes \mathbb{C}_{base} , which can be formally represented as:

$$\mathbf{B} = \left[\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(|\mathbb{C}_{\text{base}}|)} \right]. \quad (8)$$

And $\mathbf{b}^{(i)}$ is the feature vector of base class i , which is defined as:

$$\mathbf{b}^{(i)} = \frac{1}{M^{(i)}} \sum_{j=1}^{M^{(i)}} \text{GAP}(\mathbf{F}^{(i,j)}), \quad (9)$$

where $M^{(i)}$ is the number of samples of base class i . $\text{GAP}(\cdot)$ is global average pooling defined in Eq. (5). $\mathbf{F}^{(i,j)}$ is the feature map of the j -th sample of base class i in \mathbb{D}_{base} .

As can be seen from the above, Cooperative Bi-path Metric is a nonparametric (model-free) method, if we do not consider the selection of similarity functions σ , σ' and weight hyperparameter α . It does not introduce additional network parameters or change the training process, only additionally takes the similarity distributions of the support samples and the query samples on base classes into consideration. We can just simply append Cooperative Bi-path Metric to any trained models. However, in this way, the classification of query samples depends not only on a small number of support samples but also on the information provided by base classes, thus increasing the robustness of the model when support samples are insufficient.

3.4 Revisiting Few-shot Learning with LLE

According to Eq. (6) and (7) in Section 3.3, each base class makes equal contribution to $\varphi^{(n)}$, while ρ_{query} and $\rho_{\text{support}}^{(n)}$ are linear about all base classes without focusing on some specific classes. Thus there arises a concern: for each query sample, different base vectors should contribute differently based on the correlations in the latent space. For example, the *walker foxhound* from base classes should be prominent when querying the *golden retriever* sample.

Cooperative Bi-Path metric with LLE. We replace ρ_{query} and $\rho_{\text{support}}^{(n)}$ with nonlinear $\tilde{\rho}_{\text{query}}$ and $\tilde{\rho}_{\text{support}}^{(n)}$ through using local linear embedding (LLE) [40]. Compared with the conventional dimensionality reduction methods such as PCA and LDA which focus on sample variance, LLE focuses on maintaining the local linear characteristics of samples when reducing sample dimensionality. LLE assumes that each sample can be represented by linearly combining its k nearest neighbors, and the weight coefficient of the linear relationship before and after dimensionality reduction remains unchanged. It can be seen that LLE has some selectivity in the dimensionality reduction process, which meets our expectation that the samples should focus on some specific base classes. The process of Cooperative Bi-path Metric with LLE is shown in Alg. 1.

Algorithm 1: Cooperative Bi-path Metric with LLE

Input: $\mathbf{B} = [\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(|\mathbb{C}_{base}|)}] \in \mathbb{R}^{c \times |\mathbb{C}_{base}|}$: feature matrix of base classes, $\mathbf{q} \in \mathbb{R}^c$: feature vector of query sample, $\mathbf{s}^{(n)} \in \mathbb{R}^c$: feature vector of novel class n , k : number of nearest neighbors to be considered for each sample, c' : dimensionality after reduction, α : weight hyperparameter introduced in Section 3.3

Output: $\tilde{\psi}^{(n)}$: final classification score of the query sample for novel class n

- 1 **for** i in $1, \dots, |\mathbb{C}_{base}|$ **do**
- 2 Find the k nearest neighbors of $\mathbf{b}^{(i)}$ in Euclidean distance:
 $\mathbf{N}^{(i)} = KNN(\mathbf{b}^{(i)}, k) = [\mathbf{b}^{(i,1)}, \dots, \mathbf{b}^{(i,k)}] \in \mathbb{R}^{c \times k}$;
- 3 Calculate the local covariance matrix:
 $\mathbf{C}^{(i)} = (\mathbf{B}^{(i)} - \mathbf{N}^{(i)})^\top (\mathbf{B}^{(i)} - \mathbf{N}^{(i)}) \in \mathbb{R}^{k \times k}$, where
 $\mathbf{B}^{(i)} = [\mathbf{b}^{(i)}, \dots, \mathbf{b}^{(i)}] \in \mathbb{R}^{c \times k}$;
- 4 Calculate the weight coefficient vector:
 $\mathbf{w}^{(i)} = \frac{\mathbf{C}^{(i)-1} \mathbf{1}_k}{\mathbf{1}_k^\top \mathbf{C}^{(i)-1} \mathbf{1}_k} \in \mathbb{R}^k$, where $\mathbf{1}_k \in \mathbb{R}^k$ is a vector full of 1 and $\mathbf{C}^{(i)-1}$ is the inverse matrix of $\mathbf{C}^{(i)}$;
- 5 **end**
- 6 Calculate matrix:
 $\mathbf{M} = (\mathbf{I}_{|\mathbb{C}_{base}|} - \mathbf{W}) (\mathbf{I}_{|\mathbb{C}_{base}|} - \mathbf{W})^\top \in \mathbb{R}^{|\mathbb{C}_{base}| \times |\mathbb{C}_{base}|}$, where
 $\mathbf{I}_{|\mathbb{C}_{base}|} \in \mathbb{R}^{|\mathbb{C}_{base}| \times |\mathbb{C}_{base}|}$ is a identity matrix and \mathbf{W} 's each element
 $\mathbf{W}_{j,i} = \begin{cases} \mathbf{w}_k^{(i)} & \text{if } \mathbf{b}^{(j)} \text{ is the } k\text{-th neighbor of } \mathbf{b}^{(i)}, \\ 0 & \text{otherwise.} \end{cases}$
- 7 Obtain dimensionality reduced feature matrix:
 $\tilde{\mathbf{B}} = [\tilde{\mathbf{b}}^{(1)}, \dots, \tilde{\mathbf{b}}^{(|\mathbb{C}_{base}|)}] \in \mathbb{R}^{c' \times |\mathbb{C}_{base}|}$, where i -th row vector in $\tilde{\mathbf{B}}$ is the eigenvector corresponding to the $(i + 1)$ -th smallest eigenvalue of the matrix \mathbf{M} ;
- 8 **for** \mathbf{q} **do**
- 9 Find the k nearest neighbors of \mathbf{q} in \mathbf{B} as step 2:
 $\mathbf{N}_{query} = [\mathbf{b}_{query}^{(1)}, \dots, \mathbf{b}_{query}^{(k)}] \in \mathbb{R}^{c \times k}$;
- 10 Find k corresponding dimensionality reduced vectors of \mathbf{N}_{query} in $\tilde{\mathbf{B}}$: $\tilde{\mathbf{N}}_{query} = [\tilde{\mathbf{b}}_{query}^{(1)}, \dots, \tilde{\mathbf{b}}_{query}^{(k)}] \in \mathbb{R}^{c' \times k}$;
- 11 Calculate the weight coefficient vector $\mathbf{w}_{query} \in \mathbb{R}^k$ as step 3 and 4;
- 12 Calculate the dimensionality reduced feature vector of query sample: $\tilde{\mathbf{q}} = \tilde{\mathbf{N}}_{query} \mathbf{w}_{query} \in \mathbb{R}^{c'}$;
- 13 **end**
- 14 **for** $\mathbf{s}^{(n)}$ **do**
- 15 Calculate the dimensionality reduced feature vector $\tilde{\mathbf{s}}^{(n)} \in \mathbb{R}^{c'}$ of novel class n as step 9-12;
- 16 **end**
- 17 Calculate nonlinear $\tilde{\rho}_{query}$, $\tilde{\rho}_{support}^{(n)}$ and $\tilde{\psi}^{(n)}$ according to Eq. (6) and (7) with dimensionality reduced $\tilde{\mathbf{B}}$, $\tilde{\mathbf{q}}$, $\tilde{\mathbf{s}}^{(n)}$ and weight hyperparameter α ;

In Alg. 1, once $\tilde{\rho}_{query}$, $\tilde{\rho}_{support}^{(n)}$ and $\tilde{\psi}^{(n)}$ are obtained, query samples are assigned to novel class \hat{y} with maximum classification score. Besides, Cooperative Bi-path Metric with LLE increases non-linearity between base classes and novel samples through reducing dimensionality with LLE, thus different base classes can make different influence on classifying different query samples by the process of finding their k nearest neighbors.

4 EXPERIMENTS

4.1 Experiment Setting

Datasets. We conduct experiments on two widely-used benchmarks, *i.e.*, miniImageNet [4] and tieredImageNet [12]. MiniImageNet is a subset of ImageNet [1], with 100 classes and 600 images in each class. Among these classes, 64 are for training, 16 for validation, and 20 for testing. TieredImageNet is another much larger subset of ImageNet, with 34 categories (608 classes) and 779,165 images in total. There are 20 categories (351 classes) for training, 6 categories (97 classes) for validation, and 8 categories (160 classes) for testing. Different from miniImageNet, tieredImageNet has no semantic overlap between the training set, test set, and validation set, so it is a potentially harder few-shot learning benchmark. All images are resized to 84×84 pixels.

Training. We follow the same experiment setting details as previous works, *e.g.* [17]. For convenience, we still organize the training data in the *episodic* format for all methods, even if some methods use only query samples in the way of supervised learning. For the setting of N -way K -shot, each task samples N classes from the novel classes \mathbb{C}_{base} , and further samples K support samples and 6 query samples for each class. For all models, we train them for 90 epochs on miniImageNet and 80 epochs on tieredImageNet, and each epoch contains 1,200 tasks.

Validation and evaluation. After each epoch, we save the best models with maximum classification accuracy and tune the hyperparameters according to the validation set, and finally report the models' accuracy on the testing set. Similar to the training phase, the data are also organized into tasks during validation and testing phases, except that 15 query samples are sampled for each class. There are 2,000 tasks during the validation and testing phase, and we reported the models' accuracy and corresponding 95% confidence interval on these tasks.

Implementation details. We implement our model using Pytorch [41], with the part of LLE using scikit-learn [42]. All experiments are performed on a single consumer-level NVIDIA 2080Ti GPU. As mentioned in Section 3.2, we use ResNet-12 as the backbone and adopt SGD with Nesterov momentum of 0.9 and weight-decay of 5×10^{-4} to optimize the models. The initial learning rate is set as 0.1, decreases to 0.006 at 60 epochs, and decreases to 20% per 10 epochs thereafter. Each mini-batch contains 4 tasks during training. Except otherwise stated, the temperature hyperparameter t is fixed as 0.6 for miniImageNet and 0.7 for tieredImageNet. For Cooperative Bi-path Metric, the hyperparameter values, α , k , and c' , all take the values which produce the highest accuracy on the validation set under different settings. The code and model are at <http://cvteam.net/projects/2020/CBM>.

Table 1: Comparison to prior works on 5-way classification on miniImageNet benchmark. ConvNet is a 4-layer convolutional network, WideResNet is a wider version of ResNet, baseline++ denotes our baseline with tricks and CBM (CBM+LLE) indicates Cooperative Bi-path Metric (with LLE). All the numbers of prior works are imported from corresponding original papers. The best results are bolded.

Model	Backbone	1-shot	5-shot
MAML [7]	ConvNet	48.70 ± 0.84	55.31 ± 0.73
MN [4]	ConvNet	46.6	60.0
ML-LSTM [6]	ConvNet	43.44 ± 0.77	60.60 ± 0.71
mAP-SSVM [8]	ConvNet	50.32 ± 0.80	63.94 ± 0.72
Meta-SGD [5]	ConvNet	50.47 ± 1.87	64.03 ± 0.94
Ren <i>et al.</i> [12]	ConvNet	50.41 ± 0.31	64.39 ± 0.24
RN [11]	ConvNet	50.44 ± 0.82	65.32 ± 0.70
Reptile [10]	ConvNet	49.97 ± 0.32	65.99 ± 0.58
PN [9]	ConvNet	49.92 ± 0.78	68.20 ± 0.66
DN4 [15]	ConvNet	51.24 ± 0.74	71.02 ± 0.64
Gidaris <i>et al.</i> [18]	ConvNet	56.20 ± 0.86	73.00 ± 0.64
Qiao <i>et al.</i> [19]	WideResNet	59.60 ± 0.41	73.74 ± 0.19
SNAIL [13]	ResNet-12	55.71 ± 0.99	68.88 ± 0.92
TEAM [16]	ResNet-12	60.07	75.90
TADAM [14]	ResNet-12	58.50 ± 0.30	76.70 ± 0.30
CAN [17]	ResNet-12	63.85 ± 0.48	79.44 ± 0.34
DC [20]	ResNet-12	62.53 ± 0.19	79.77 ± 0.19
baseline++ (ours)	ResNet-12	64.07 ± 0.45	80.47 ± 0.33
CBM (ours)	ResNet-12	64.77 ± 0.46	80.50 ± 0.33
CBM+LLE (ours)	ResNet-12	64.21 ± 0.45	80.68 ± 0.32

4.2 Comparison with the State of the Art

To evaluate the effectiveness of the proposed approach, we first conduct experiments on the miniImageNet [4] benchmark. As shown in Tab. 1, it can be found that although the baseline is a simple network with a fully-connected layer trained in the way of traditional supervised learning, it has achieved comparable results to the state of the art. This suggests that episodic training procedure is unnecessary and makes us rethink whether the previous complex episodic methods make sense. And we recommend this baseline for future study.

Also, Cooperative Bi-path Metric outperforms the baseline by a relatively large margin in the 5-way 1-shot setting, with an improvement of about 0.9% compared to the state of the art. This indicates that the base classes help classify query samples of novel classes, especially in cases where the labeled support samples are severely insufficient.

Cooperative Bi-path Metric with LLE can slightly improve the accuracy compared with the vanilla version in the 5-way 5-shot setting, and it is about 0.9% higher than the state of the art, which illustrates the importance of how to adaptively utilize different base classes for classifying different samples of novel classes. However, the information of base classes is underused in existing methods, we regard this as a promising but underappreciated direction in few-shot learning.

Table 2: Comparison to prior works on 5-way classification on tieredImageNet benchmark. The best results are bolded.

Model	Backbone	1-shot	5-shot
MAML [7]	ConvNet	51.67 ± 1.81	70.30 ± 1.75
Ren <i>et al.</i> [12]	ConvNet	52.39 ± 0.44	69.88 ± 0.20
RN [11]	ConvNet	54.48 ± 0.93	71.32 ± 0.78
PN [9]	ConvNet	53.31 ± 0.89	72.69 ± 0.74
CAN [17]	ResNet-12	69.89 ± 0.51	84.23 ± 0.37
baseline++ (ours)	ResNet-12	71.27 ± 0.50	85.81 ± 0.34

Table 3: Results of different training modes on 5-way classification on miniImageNet. GL: global loss; FL: few-shot loss; GL+FL: the average of global loss and few-shot loss. The best results are bolded.

Model	Loss type	1-shot	5-shot
TADAM [14]	GL	52.60 ± 0.92	75.10 ± 0.67
	FL	52.74 ± 0.91	73.42 ± 0.71
	GL+FL	55.97 ± 0.93	73.48 ± 0.68
baseline++	GL	64.07 ± 0.45	80.47 ± 0.33
	FL	58.05 ± 0.49	74.52 ± 0.37
	GL+FL	63.15 ± 0.47	77.65 ± 0.35

It is worth noting that Cooperative Bi-path Metric classifies query samples based on the same trained backbone of the baseline, and the accuracy improvement over the baseline is stable and not subject to the randomness of the training procedure. Compared to the baseline, Cooperative Bi-path Metric does not introduce additional network parameters or model updating processes, just changes the classification criterion during the testing phase. It is a computationally lightweight method, can be easily integrated into other trained models, such as the Prototypical Networks [9] or Matching Networks [4].

To further illustrate the effectiveness of traditional supervised learning, we also compare our baseline on tieredImageNet [12] with existing methods. As we can see from Tab. 2, our baseline also gets the best results in both settings. Notably, the baseline and the proposed Cooperative Bi-path Metric (as well as the version with LLE) set up a new state of the art in the field of few-shot learning.

4.3 Rethinking Few-shot Training Mode

Is episodic training necessary for few-shot learning? To make the training mode and testing mode consistent as well as to get better performance, the previous few-shot learning methods adopted the *episodic* training process. More specifically, similar to the test phase, they also sampled many N -way K -shot tasks during the training phase to train a model with the cross-entropy loss of query samples over N classes in each task. We call this kind of loss *few-shot loss*. Different from the episodic training mode with few-shot loss, the other training mode does not only focus on the classes within a task but adopts the traditional supervised learning, using the whole training set \mathbb{D}_{base} to train a feature extractor and a

Table 4: Influence of the three tricks on the baseline on 5-way classification on miniImageNet benchmark. DA: data augmentation; t : temperature; DC: dense classification.

DA	t	DC	1-shot	5-shot
			51.68 ± 0.44	69.53 ± 0.36
✓			59.72 ± 0.44	77.67 ± 0.34
	✓		52.09 ± 0.44	69.80 ± 0.37
		✓	54.00 ± 0.45	70.74 ± 0.36
✓	✓	✓	64.07 ± 0.45	80.47 ± 0.33

$|\mathbb{C}_{base}|$ -way fully-connected layer. In this mode, the entire network is trained with the cross-entropy loss of samples over all classes in the training set. We call this kind of loss *global loss*. To study the impact of different training modes on the model’s accuracy, we train TADAM [14] and our proposed baseline in different modes. For convenience, in the experiments involving global loss, we also organize training data in the form of tasks, but only calculate query samples’ global loss. For both modes, during the testing phase, the feature extractor is used to extract the features of all the samples within a task, and query samples are classified into the class with the maximum inductive similarity.

It can be seen from Tab. 3 that for TADAM [14], global loss alone is similar to few-shot loss alone in 1-shot setting, but it is much better than few-shot loss in 5-shot. And it is interesting to find that using both global loss and few-shot loss works worse than global loss alone in 5-shot, *i.e.*, the introduction of few-shot loss reduces the performance of global loss. For baseline++, we can also find that global loss is better than few-shot loss in both settings, and the gap between themselves is larger than that for TADAM [14]. Through this experiment, we realize that the episodic training mode is not necessary and its capacity is limited. We also believe that using as much global information as possible to train an efficient feature extractor is important for few-shot learning.

During both training and testing phases, TADAM [14] uses Euclidean distance to measure the distance between samples and classes, while baseline++ uses cosine similarity. We follow the same training strategy as TADAM [14], but the produced results are slightly different from the reported ones.

4.4 Bag of Tricks for Strong Baseline

To explore the influence of the three tricks (data enhancement, temperature, and dense classification) on the baseline, we conduct ablation experiments on miniImageNet. The experimental results are shown in Tab. 4, from which we can find that these tricks can greatly improve the accuracy of the baseline, especially data enhancement.

To further explore the impact of different values of temperature t on the baseline’s accuracy, we conduct experiments with different values of t , and the corresponding results are shown in Fig. 3. It can be seen from Fig. 3 that the accuracy curves are upward convex, and the highest accuracies are obtained when t is 0.6 in both settings. All the models proposed in this paper adopt 0.6 as the value of t on miniImageNet and 0.7 on tieredImageNet.

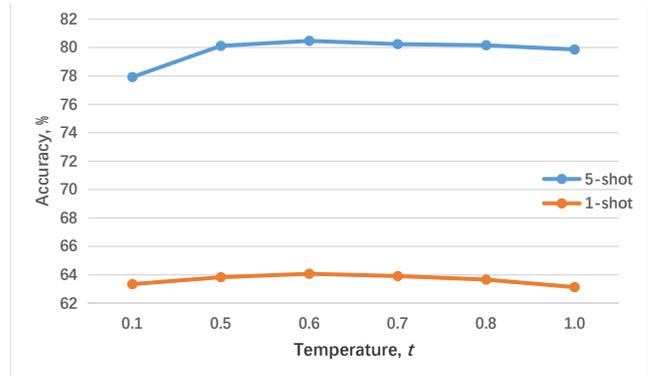


Figure 3: The accuracy curves of baseline for different values of t on 5-way classification on miniImageNet benchmark. All experiments contain both data augmentation and dense classification.

Table 5: Accuracy of the variants of vanilla Cooperative Bi-path Metric on 5-way classification on miniImageNet benchmark. CS: cosine similarity; ED: Euclidean distance; KL: KL divergence. The best results are bolded.

σ'	softmax	σ	1-shot		5-shot	
			α	Acc.	α	Acc.
CS	No	CS	0.15	64.60 ± 0.46	0.75	80.48 ± 0.33
		ED	0.80	64.52 ± 0.46	1.00	80.47 ± 0.33
	Yes	CS	0.05	64.77 ± 0.46	0.35	80.50 ± 0.33
		ED	0.05	64.75 ± 0.46	0.65	80.49 ± 0.33
ED	No	CS	0.20	64.62 ± 0.46	0.85	80.48 ± 0.33
		ED	0.95	64.20 ± 0.45	1.00	80.47 ± 0.33
	Yes	CS	0.10	64.43 ± 0.45	0.70	80.48 ± 0.33
		ED	0.10	64.39 ± 0.45	0.80	80.48 ± 0.33
		ED	0.10	64.45 ± 0.45	0.70	80.49 ± 0.33
		KL	0.10	64.45 ± 0.45	0.70	80.49 ± 0.33

4.5 Variants of Cooperative Bi-path Metric

Variants of vanilla Cooperative Bi-path Metric. As we can see from the Tab. 1, Cooperative Bi-path Metric, as well as the version with LLE, can further improve classification accuracy by utilizing base classes during testing phase. However, we found that the specific implementation of different details of Cooperative Bi-path Metric has a considerable impact on the performance, and different variants of Cooperative Bi-path Metric have different accuracy. For the vanilla Cooperative Bi-path Metric, it needs to consider the specific implementation of these five details: (i) should cosine similarity or Euclidean distance be used as the similarity function σ' to calculate ρ_{query} and $\rho_{support}^{(n)}$ for query samples and support samples? (ii) after obtaining ρ_{query} and $\rho_{support}^{(n)}$, whether softmax is applied on ρ_{query} and $\rho_{support}^{(n)}$ to obtain normalized similarity distribution (all components add up to 1)? (iii) should cosine similarity

Table 6: Accuracy of the variants of Cooperative Bi-path Metric with LLE on 5-way classification on miniImageNet benchmark. The dimensionality after reduction c' is fixed as 63. The best results are bolded.

L_2	σ'	softmax	σ	1-shot			5-shot		
				k	α	Acc.	k	α	Acc.
No	No	CS	8	0.95	64.16 ± 0.45	24	0.95	80.68 ± 0.32	
			ED	8	0.95	64.17 ± 0.45	24	0.95	80.67 ± 0.32
	Yes	CS	10	0.35	64.21 ± 0.45	26	0.35	80.58 ± 0.32	
			ED	11	0.30	64.19 ± 0.45	24	0.30	80.59 ± 0.32
	No	KL	10	0.35	64.20 ± 0.45	25	0.25	80.63 ± 0.32	
				0.25	64.14 ± 0.45	23	0.10	80.55 ± 0.32	
ED	No	CS	4	0.25	64.14 ± 0.45	23	0.10	80.55 ± 0.32	
			ED	5	1.00	64.07 ± 0.45	63	0.95	80.49 ± 0.32
	Yes	CS	21	0.55	64.11 ± 0.45	33	0.60	80.50 ± 0.32	
			ED	22	0.55	64.11 ± 0.45	23	0.70	80.51 ± 0.32
	No	KL	22	0.50	64.11 ± 0.45	23	0.40	80.51 ± 0.32	
				0.40	64.07 ± 0.45	63	0.95	80.49 ± 0.32	
CS	No	CS	7	0.95	64.20 ± 0.45	23	0.95	80.64 ± 0.32	
			ED	8	0.95	64.20 ± 0.45	22	0.95	80.63 ± 0.32
	Yes	CS	5	0.30	64.19 ± 0.45	12	0.55	80.57 ± 0.32	
			ED	9	0.40	64.18 ± 0.45	25	0.30	80.58 ± 0.32
	No	KL	7	0.30	64.20 ± 0.45	23	0.30	80.61 ± 0.32	
				0.20	64.18 ± 0.45	23	0.30	80.61 ± 0.32	
ED	No	CS	16	1.00	64.16 ± 0.45	25	0.10	80.58 ± 0.32	
			ED	5	1.00	64.07 ± 0.45	51	0.95	80.65 ± 0.32
	Yes	CS	5	0.45	64.12 ± 0.45	23	0.75	80.51 ± 0.32	
			ED	5	0.40	64.12 ± 0.45	20	0.65	80.51 ± 0.32
	No	KL	7	0.30	64.12 ± 0.45	22	0.70	80.51 ± 0.32	
				0.20	64.12 ± 0.45	22	0.70	80.51 ± 0.32	

or Euclidean distance or KL divergence be used as the similarity function σ to calculate the transductive similarity $\phi^{(n)}$ between ρ_{query} and $\rho_{support}^{(n)}$? (iv) how to balance the inductive similarity $\phi^{(n)}$ and the transductive similarity $\phi^{(n)}$ by weight hyperparameter α ? For (i) - (iii), we enumerate and experiment with all possible combinations. For each combination, the highest accuracy is found through changing α in range $[0, 1]$ with a interval of 0.05. All results are shown in Tab. 5. It can be found that under the setting of 1-shot, Cooperative Bi-path Metric can greatly improve the accuracy (compared with 64.07), and the value of α is quite small, *i.e.*, the transductive similarity $\phi^{(n)}$ plays a major role. However, the result under 5-shot is not ideal. The improvement of accuracy is limited (compared with 80.47), and the value of α is large, so inductive similarity $\phi^{(n)}$ is still in the dominant place.

Variants of Cooperative Bi-path Metric with LLE. For Cooperative Bi-path Metric with LLE, it requires additional determination of (v) whether or not to conduct L_2 normalization on feature vectors before LLE and the values of two additional hyperparameters, namely (vi) the number of nearest neighbors k and (vii) the dimensionality after reduction c' in LLE. Like vanilla Cooperative Bi-path Metric, for (i) - (vi), we enumerate and experiment with all possible combinations. For (vii), through some preliminary experiments, we find that the highest accuracy of various combinations are obtained when c' is 63, so we fix c' as 63 for all combinations. And the highest accuracy under various combinations is shown in Tab. 6.

Influence of weight hyperparameter. In order to further explore the influence of weight hyperparameter α on classification accuracy (especially when $\alpha = 0$, only transductive similarity $\phi^{(n)}$ works), we report the accuracy curves with different values of α , as shown

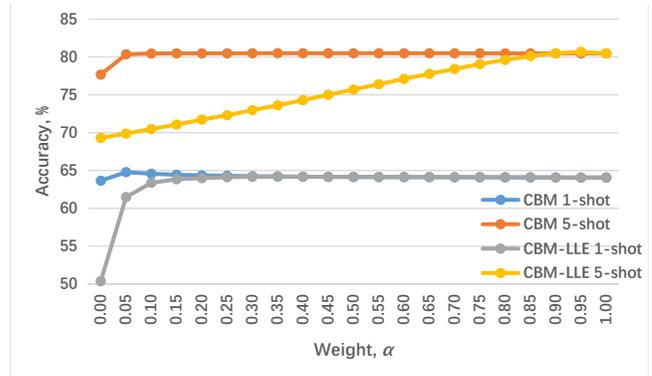


Figure 4: The accuracy curves for different values of α on 5-way classification on miniImageNet benchmark.

in Fig. 4. It can be seen from Fig. 4, when $\phi^{(n)}$ is the majority, the accuracy is higher, $\phi^{(n)}$ contributes more to the classification than $\phi^{(n)}$. We attribute part of the reason why $\phi^{(n)}$ takes a larger part to that the magnitude of $\phi^{(n)}$ and $\phi^{(n)}$ is different. The accuracy at $\alpha = 0$ is less than the accuracy at $\alpha = 1$, which means that classification result based on $\phi^{(n)}$ alone is less accurate than that based on $\phi^{(n)}$ alone. However, the highest points of the accuracy curves are not obtained at $\alpha = 1$, so $\phi^{(n)}$ can further improve the classification accuracy based on $\phi^{(n)}$.

5 CONCLUSIONS AND FUTURE WORK

In this work, we contribute to few-shot learning with a concise and effective baseline as well as a novel metric named Cooperative Bi-path Metric.

First, we train a simple network in the way of traditional supervised learning as the baseline, which achieves comparable results to the state of the art. This shows that episodic training mode is not necessary, and an effective feature extractor to capture discriminative features of samples is fundamental for few-shot learning.

Second, we propose Cooperative Bi-path Metric to change the criterion of classification. It uses samples' similarity distribution on base classes to assist classification decisions while existing methods did not take full use of such information of base classes. Experiments show that it can further boost the model's performance and achieve a new state of the art in the field of few-shot image classification, indicating that using base classes to classify samples during the testing phase looks like a promising direction for future research.

However, Cooperative Bi-path Metric is handcrafted and somewhat straightforward. A natural direction for improving it is training an additional convolutional neural network end to end to measure the transductive similarity. We leave this for future work.

ACKNOWLEDGMENTS

This work was supported in part by the National Key R&D Program of China (No. 2017YFB1002400), the National Natural Science Foundation of China (No. 61922006, No. 61532003, No. 61825101 and No. U1611461), and the Beijing Nova Program (No. Z181100006218063).

REFERENCES

- [1] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. 2015. ImageNet Large Scale Visual Recognition Challenge. *IJCV* 115, 3 (2015), 211–252. <https://doi.org/10.1007/s11263-015-0816-y>
- [2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *NeurIPS*. 1106–1114. <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks>
- [3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [4] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching Networks for One Shot Learning. In *NeurIPS*. 3630–3638. <http://papers.nips.cc/paper/6385-matching-networks-for-one-shot-learning>
- [5] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-SGD: Learning to Learn Quickly for Few Shot Learning. *CoRR* abs/1707.09835 (2017). arXiv:1707.09835 <http://arxiv.org/abs/1707.09835>
- [6] Sachin Ravi and Hugo Larochelle. 2017. Optimization as a Model for Few-Shot Learning. In *ICLR*. <https://openreview.net/forum?id=rjY0-Kcll>
- [7] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 70)*. 1126–1135. <http://proceedings.mlr.press/v70/finn17a.html>
- [8] Eleni Triantafillou, Richard S. Zemel, and Raquel Urtasun. 2017. Few-Shot Learning Through an Information Retrieval Lens. In *NeurIPS*. 2255–2265. <http://papers.nips.cc/paper/6820-few-shot-learning-through-an-information-retrieval-lens>
- [9] Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. Prototypical Networks for Few-shot Learning. In *NeurIPS*. 4077–4087. <http://papers.nips.cc/paper/6996-prototypical-networks-for-few-shot-learning>
- [10] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On First-Order Meta-Learning Algorithms. *CoRR* abs/1803.02999 (2018). arXiv:1803.02999 <http://arxiv.org/abs/1803.02999>
- [11] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. 2018. Learning to Compare: Relation Network for Few-Shot Learning. In *CVPR*. 1199–1208. <https://doi.org/10.1109/CVPR.2018.00131>
- [12] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. 2018. Meta-Learning for Semi-Supervised Few-Shot Classification. In *ICLR*. <https://openreview.net/forum?id=HJcSzz-CZ>
- [13] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. A Simple Neural Attentive Meta-Learner. In *ICLR*. <https://openreview.net/forum?id=B1DmUzWAW>
- [14] Boris N. Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. 2018. TADAM: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*. 719–729. <http://papers.nips.cc/paper/7352-tadam-task-dependent-adaptive-metric-for-improved-few-shot-learning>
- [15] Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo. 2019. Revisiting Local Descriptor Based Image-To-Class Measure for Few-Shot Learning. In *CVPR*. 7260–7268. <https://doi.org/10.1109/CVPR.2019.00743>
- [16] Limeng Qiao, Yemin Shi, Jia Li, Yonghong Tian, Tiejun Huang, and Yaowei Wang. 2019. Transductive Episodic-Wise Adaptive Metric for Few-Shot Learning. In *ICCV*. 3602–3611. <https://doi.org/10.1109/ICCV.2019.00370>
- [17] Ruibing Hou, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. 2019. Cross Attention Network for Few-shot Classification. In *NeurIPS*. 4005–4016. <http://papers.nips.cc/paper/8655-cross-attention-network-for-few-shot-classification>
- [18] Spyros Gidaris and Nikos Komodakis. 2018. Dynamic Few-Shot Visual Learning Without Forgetting. In *CVPR*. 4367–4375. <https://doi.org/10.1109/CVPR.2018.00459>
- [19] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L. Yuille. 2018. Few-Shot Image Recognition by Predicting Parameters From Activations. In *CVPR*. 7229–7238. <https://doi.org/10.1109/CVPR.2018.00755>
- [20] Yann Lifchitz, Yannis Avrithis, Sylvaine Picard, and Andrei Bursuc. 2019. Dense Classification and Implanting for Few-Shot Learning. In *CVPR*. 9258–9267. <https://doi.org/10.1109/CVPR.2019.00948>
- [21] Guneet S. Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. 2019. A Baseline for Few-Shot Image Classification. *CoRR* abs/1909.02729 (2019). arXiv:1909.02729 <http://arxiv.org/abs/1909.02729>
- [22] Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. 1997. Locally Weighted Learning. *Artif. Intell. Rev.* 11, 1-5 (1997), 11–73. <https://doi.org/10.1023/A:1006559212014>
- [23] Jacob Goldberger, Sam T. Roweis, Geoffrey E. Hinton, and Ruslan Salakhutdinov. 2004. Neighbourhood Components Analysis. In *NeurIPS*. 513–520. <http://papers.nips.cc/paper/2566-neighbourhood-components-analysis>
- [24] Michael Fink. 2004. Object Classification from a Single Example Utilizing Class Relevance Metrics. In *NeurIPS*. 449–456. <http://papers.nips.cc/paper/2576-object-classification-from-a-single-example-utilizing-class-relevance-metrics>
- [25] Sumit Chopra, Raia Hadsell, and Yann LeCun. 2005. Learning a Similarity Metric Discriminatively, with Application to Face Verification. In *CVPR*. 539–546. <https://doi.org/10.1109/CVPR.2005.202>
- [26] Kilian Q. Weinberger, John Blitzer, and Lawrence K. Saul. 2005. Distance Metric Learning for Large Margin Nearest Neighbor Classification. In *NeurIPS*. 1473–1480. <http://papers.nips.cc/paper/2795-distance-metric-learning-for-large-margin-nearest-neighbor-classification>
- [27] Ruslan Salakhutdinov and Geoffrey E. Hinton. 2007. Learning a Nonlinear Embedding by Preserving Class Neighbourhood Structure. *JMLR* 2, 412–419. <http://proceedings.mlr.press/v2/salakhutdinov07a.html>
- [28] Kilian Q. Weinberger and Lawrence K. Saul. 2009. Distance Metric Learning for Large Margin Nearest Neighbor Classification. *JMLR* 10 (2009), 207–244. <https://dl.acm.org/citation.cfm?id=1577078>
- [29] Martin Köstinger, Martin Hirzer, Paul Wohlhart, Peter M. Roth, and Horst Bischof. 2012. Large scale metric learning from equivalence constraints. In *CVPR*. 2288–2295. <https://doi.org/10.1109/CVPR.2012.6247939>
- [30] Aurélien Bellet, Amaury Habrard, and Marc Sebban. 2013. A Survey on Metric Learning for Feature Vectors and Structured Data. *CoRR* abs/1306.6709 (2013). arXiv:1306.6709 <http://arxiv.org/abs/1306.6709>
- [31] Brian Kulis. 2013. Metric Learning: A Survey. *Foundations and Trends in Machine Learning* 5, 4 (2013), 287–364. <https://doi.org/10.1561/22000000019>
- [32] Elad Hoffer and Nir Ailon. 2015. Deep metric learning using Triplet network. In *ICLR*. <http://arxiv.org/abs/1412.6622>
- [33] Mehrtash Tafazzoli Harandi, Mathieu Salzmann, and Richard I. Hartley. 2017. Joint Dimensionality Reduction and Metric Learning: A Geometric Take. In *ICML (Proceedings of Machine Learning Research, Vol. 70)*. 1404–1413. <http://proceedings.mlr.press/v70/harandi17a.html>
- [34] Sebastian Thrun. 1998. Lifelong Learning Algorithms. In *Learning to Learn*, Sebastian Thrun and Lorien Y. Pratt (Eds.). Springer, 181–209. https://doi.org/10.1007/978-1-4615-5529-2_8
- [35] Sebastian Thrun and Lorien Y. Pratt. 1998. Learning to Learn: Introduction and Overview. In *Learning to Learn*, Sebastian Thrun and Lorien Y. Pratt (Eds.). Springer, 3–17. https://doi.org/10.1007/978-1-4615-5529-2_1
- [36] Ricardo Vilalta and Youssef Drissi. 2002. A Perspective View and Survey of Meta-Learning. *Artif. Intell. Rev.* 18, 2 (2002), 77–95. <https://doi.org/10.1023/A:1019956318069>
- [37] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- [38] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. 2017. Random Erasing Data Augmentation. *CoRR* abs/1708.04896 (2017). arXiv:1708.04896 <http://arxiv.org/abs/1708.04896>
- [39] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. *CoRR* abs/1503.02531 (2015). arXiv:1503.02531 <http://arxiv.org/abs/1503.02531>
- [40] Sam T Roweis and Lawrence K Saul. 2000. Nonlinear dimensionality reduction by locally linear embedding. *science* 290, 5500 (2000), 2323–2326.
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*. 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library>
- [42] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *JMLR* 12 (2011), 2825–2830.