

Few-Shot Classification with Contrastive Learning

Zhanyuan Yang¹, Jinghua Wang², and Yingying Zhu^{*1}

¹ College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China

yangzhanyuan2019@email.szu.edu.cn, zhuyy@szu.edu.cn

² School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen), Shenzhen, China
wangjinghua@hit.edu.cn

Abstract. A two-stage training paradigm consisting of sequential pre-training and meta-training stages has been widely used in current few-shot learning (FSL) research. Many of these methods use self-supervised learning and contrastive learning to achieve new state-of-the-art results. However, the potential of contrastive learning in both stages of FSL training paradigm is still not fully exploited. In this paper, we propose a novel contrastive learning-based framework that seamlessly integrates contrastive learning into both stages to improve the performance of few-shot classification. In the pre-training stage, we propose a self-supervised contrastive loss in the forms of feature vector *vs.* feature map and feature map *vs.* feature map, which uses global and local information to learn good initial representations. In the meta-training stage, we propose a cross-view episodic training mechanism to perform the nearest centroid classification on two different views of the same episode and adopt a distance-scaled contrastive loss based on them. These two strategies force the model to overcome the bias between views and promote the transferability of representations. Extensive experiments on three benchmark datasets demonstrate that our method achieves competitive results.

Keywords: Few-shot learning · Meta learning · Contrastive learning · Cross-view episodic training

1 Introduction

Thanks to the availability of a large amount of annotated data, deep convolutional neural networks (CNN) [15,21,39] yield impressive results on various visual recognition tasks. However, the time-consuming and costly collection process makes it a challenge for these deep learning-based methods to generalize in real-life scenarios with scarce annotated data. Inspired by the capability of human to learn new concepts from a few examples, few-shot learning (FSL) is considered as a promising alternative to meet the challenge, as it can adapt knowledge learned from a few samples of base classes to novel tasks.

* Corresponding author.

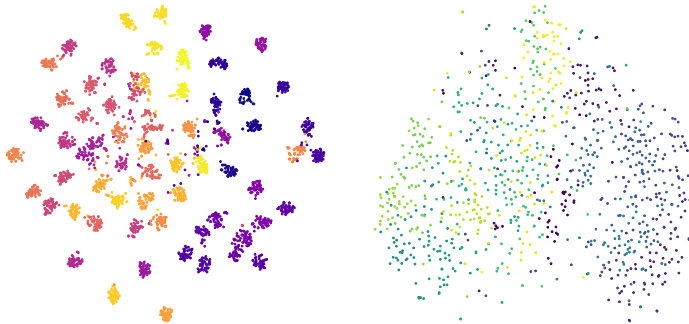


Fig. 1: Distribution of feature embeddings of 64 base (*left*) and 20 novel (*right*) classes from miniImagenet in pre-train space by t-SNE [27].

Recently, popular FSL methods [10,28,31,36,40,42,47] mainly adopt the meta-learning strategy. These meta-learning based methods typically take episodic training mechanism to perform meta-training on base classes with abundant data. During meta-training, the episodes consist of a support set and a query set, which are used in few-shot classification to mimic the evaluation setting. The learned model is expected to be capable of generalizing across novel tasks of FSL. Besides, many other methods [5,25,45,48,49,54] achieve good classification accuracy by pre-training the feature extractor on base classes. These methods suggest that the transferable and discriminative representations learned through pre-training or meta-training is crucial for few-shot classification.

However, both the pre-training and meta-training procedures only minimize the standard cross-entropy (CE) loss with labels from base classes. The resulting models are optimized to solve the classification tasks of base classes. Due to this, these methods may discard the information that might benefit the classification tasks on the unseen classes. Fig. 1 shows that the pre-trained model is able to identify samples from the base classes (left) well but performs poorly on samples from the novel classes (right). That is, the learned representations are somewhat overfitted on the base classes and not generalizable on the novel classes. Owing to the label-free nature of self-supervised learning methods, some recent works [7,11,41] have tried self-supervised pretext tasks to solve the FSL problem, while other works [9,25,26,32] focus on contrastive learning methods. Though promising, these approaches ignore the additional information from the self-supervised pretext tasks in meta-training or treat them as auxiliary losses simply in the FSL training paradigm.

In this work, we propose a contrastive learning-based framework that seamlessly integrates contrastive learning into the pre-training and meta-training stages to tackle the FSL problem. First, in the pre-training stage, we propose two types of contrastive losses based on self-supervised and supervised signals, respectively, to train the model. These losses consider the global and local in-

formation simultaneously. Our proposed self-supervised contrastive loss exploits local information in the forms of both feature vector *vs.* feature map (vector-map) and feature map *vs.* feature map (map-map), which differs from previous methods. Our supervised contrastive loss makes good use of the correlations among individual instances and the correlations among different instances of the same category. Second, in the meta-training stage, motivated by the idea of maximizing mutual information between features extracted from multiple views (*e.g.*, by applying different data augmentations on images) of the shared context (*e.g.*, original images) [1,16,30,50], we introduce a cross-view episodic training (CVET) mechanism to extract generalizable representations. Concretely, we randomly employ two different data augmentation strategies [3,14,54] to obtain the augmented episodes and treat them as different views of the original one. Note, the augmentation does not change the label of the data. We then conduct the nearest centroid classification between the augmented episodes to force the model to overcome the bias between views and generalize well to novel classes. As a complement to CVET, we take inter-instance distance scaling into consideration and perform query instance discrimination within the augmented episodes. These two methods effectively apply contrastive learning to the meta-training stage of FSL. Our proposed method learns meta-knowledge that can play a crucial role in recognizing novel classes. The key contributions of this work are as follows:

- We propose a contrastive learning-based FSL framework consisting of the pre-training and meta-training stages to improve the few-shot image classification. Our framework is easy to combine with other two-stage FSL methods.
- We adopt the self-supervised contrastive loss based on global and local information in the pre-training stage to enhance the generalizability of the resulting representations.
- We propose a CVET mechanism to force the model to find more transferable representations by executing classification between augmented episodes. Meanwhile, we introduce a distance-scaled contrastive loss based on the augmented episodes to ensure that the classification procedure is not affected by extreme bias between different views.
- Extensive experiments of few-shot classification on three benchmarks show that our proposed method achieves competitive results.

2 Related Work

2.1 Few-Shot Learning

FSL aims to learn patterns on a large number of labeled examples called base classes and adapt to novel classes with limited examples per class. Few-shot image classification has received great attention and many methods have been proposed. The existing methods can be broadly divided into two categories: *optimization-based* and *metric-based*. The *optimization-based* methods initialize

the model on base classes and adapt to novel tasks efficiently within a few gradient update steps on a few labeled samples [10,24,28,33,36]. The *metric-based* methods aim to learn a generalizable representation space and use a well-defined metric to classify them [23,31,40,42,47,51,55]. The existing works have considered different metrics such as cosine similarity [47], Euclidean distance [40], a CNN-based relation module [42], a task-adaptive metric [31], a local descriptor based metric [23] and graph neural networks [37]. The Earth Mover’s Distance [55] is employed as a metric to learn more discriminative structured representations. Many recent studies [5,38,45,48] have proposed a standard end-to-end pre-training framework to obtain feature extractors or classifiers on base classes. These pre-training based methods achieve competitive performance compared to episodic meta-training methods. Moreover, many papers [6,25,54,55] take advantage of a sequential combination of pre-training and meta-training stages to further enhance the performance. The methods [17,42,49,52,54] pay more attention to the transferability of representations through delicately designing task-specific modules in meta-training. Given the simplicity and effectiveness of these methods, we take FEAT [54] as our baseline, but drop its auxiliary loss.

2.2 Contrastive Learning

Recently, contrastive learning with the instance discrimination as a pretext task has become a dominant approach in self-supervised representation learning [1,3,14,16,30,43,50]. These methods typically construct contrast pairs of instances with a variety of data augmentations and optimize a contrastive loss with the aim of keeping instances close to their augmented counterparts while staying away from other instances in the embedding space. The goal of contrastive learning using self-supervision from instances is to improve the generalizability of the representations and benefit various downstream tasks. Contrastive learning is also extended to group instances in a supervised manner [19] and achieves better performance than CE loss on standard classification tasks.

2.3 Few-Shot Learning with Contrastive Learning

In contrast to the works [7,11,41] that introduce self-supervised pretext tasks such as rotation [13] and jigsaw [29] into FSL as auxiliary losses, recent approaches [9,25,26,32] have explored contrastive learning of instance discrimination in different parts of the two-stage training pipeline of FSL. Methods [25,26,32] combine supervised contrastive loss [19] to the pre-training stage [26,32] and the meta-training stage [9,25], respectively. Unlike prior works, our proposed method boosts few-shot classification performance by seamlessly integrating instance-discriminative contrastive learning in both the pre-training and meta-training stages. In the pre-training stage, we conduct self-supervised contrastive loss in the forms of vector-map and map-map. In the meta-training stage, we combine contrastive learning with episodic training and define a distance-scaled contrastive loss to improve the transferability of the representations.

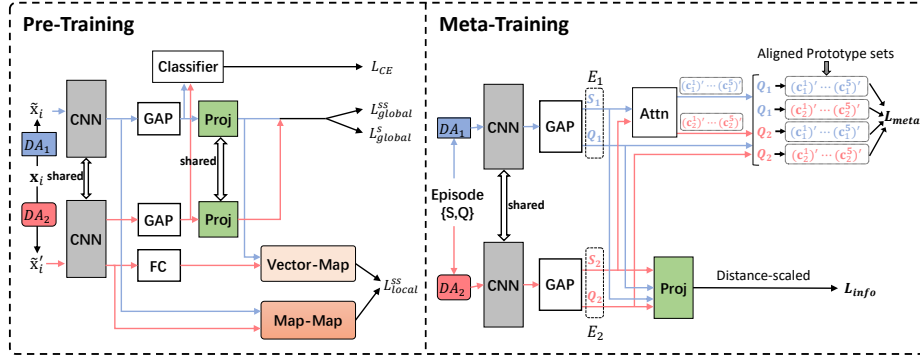


Fig. 2: Overview of our framework. Based on multiple views of an input through two random data augmentations DA_1 and DA_2 , we compute contrastive losses at both global and local levels in the pre-training stage. In the meta-training stage, we enforce cross-view episodic training and compute a distance-scaled contrastive loss episodically. Here, GAP denotes a global average pooling layer, Proj is a projection head, FC means a fully connected layer, Attn is the task-specific module from [54], \mathcal{S}_r , \mathcal{Q}_r mean support and query set from different views of the episode $E = \{\mathcal{S}, \mathcal{Q}\}$ respectively, and $(c_r^k)'$ denotes the aligned prototype in \mathcal{S}_r .

3 Method

3.1 Preliminary

The few-shot classification task is slightly different from the standard supervised classification task. The meta-training set $\mathcal{D}_{train} = \{(x_i, y_i) \mid y_i \in \mathcal{C}_{base}\}$ consists of the samples from the base classes \mathcal{C}_{base} , and the meta-test set $\mathcal{D}_{test} = \{(x_i, y_i) \mid y_i \in \mathcal{C}_{novel}\}$ consists of the samples from the novel classes \mathcal{C}_{novel} . Here, y_i is the class label of sample x_i . In FSL, we aim to learn a model based on \mathcal{D}_{train} and generalize it over \mathcal{D}_{test} , where $\mathcal{C}_{base} \cap \mathcal{C}_{novel} = \emptyset$. Following the prior meta-learning based methods [10, 40, 47], we adopt episodic mechanism to simulate the evaluation setting. Concretely, each M -way K -shot episode E consists of a support set and a query set. We first randomly sample M classes from \mathcal{C}_{base} for meta-training (or from \mathcal{C}_{novel} for meta-testing) and K instances per class to obtain the support set $\mathcal{S} = \{x_i, y_i\}_{i=1}^{M \times K}$. Then, we sample Q instances in each of the selected classes to obtain the query set $\mathcal{Q} = \{x_i, y_i\}_{i=1}^{M \times Q}$. Note that $y_i \in \{1, 2, \dots, M\}$ and $\mathcal{S} \cap \mathcal{Q} = \emptyset$. The episodic training procedure classifies the samples in \mathcal{Q} into the categories corresponding to the samples in \mathcal{S} .

3.2 Overview

In this work, we follow the two-stage training strategy and incorporate contrastive learning in both stages to learn more generalizable representations. Our

proposed framework is illustrated in Fig. 2. In the pre-training stage, we adopt self-supervised and supervised contrastive losses to obtain a good initial representation. In the meta-training stage, we propose a novel cross-view episodic training (CVET) mechanism and a distance-scaled contrastive loss, which allows the model to overcome the bias between views of each episode and generalize well across novel tasks. Note that we take FEAT [54] without its auxiliary loss as the baseline due to its simple and effective task-specific module (a multi-head attention module [46]). We will detail our framework in the following subsections.

3.3 Pre-training

In this section, we introduce instance-discriminative contrastive learning [3,19] in the pre-training stage to alleviate the overfitting problem caused by training with CE loss only. As shown in Fig. 2, we propose self-supervised contrastive losses at the global and local levels, respectively. Using self-supervision in these losses helps produce more generalizable representations. Meanwhile, we also employ a global supervised contrastive loss [19] to capture the correlations among instances from the same category.

Global self-supervised contrastive loss. This loss (*a.k.a* InfoNCE loss [3,30]) aims to enhance the similarity between the views of the same image, while reducing the similarity between the views of different images. Formally, we randomly apply two data augmentation methods to a batch of samples $\{x_i, y_i\}_{i=1}^N$ from the meta-training set \mathcal{D}_{train} and generate the augmented batch $\{\tilde{x}_i, \tilde{y}_i\}_{i=1}^{2N}$. Here, \tilde{x}_i and \tilde{x}'_i denote two different views of x_i , which are considered as a positive pair. We define f_ϕ as the feature extractor with learnable parameters ϕ to transform the sample \tilde{x}_i into a feature map $\hat{\mathbf{x}}_i = f_\phi(\tilde{x}_i) \in \mathbb{R}^{C \times H \times W}$ and further obtain the global feature $\mathbf{h}_i \in \mathbb{R}^C$ after a global average pooling (GAP) layer. We use a MLP with one hidden layer to instantiate a projection head $proj(\cdot)$ [3] to generate the projected vector $\mathbf{z}_i = proj(\mathbf{h}_i) \in \mathbb{R}^D$. Then the global self-supervised contrastive loss can be computed as:

$$L_{global}^{ss} = - \sum_{i=1}^{2N} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}'_i / \tau_1)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} \exp(\mathbf{z}_i \cdot \mathbf{z}_j / \tau_1)}, \quad (1)$$

where the \cdot operation denotes inner product after l_2 normalization, τ_1 is a scalar temperature parameter, and $\mathbb{1} \in \{0, 1\}$ is an indicator function. Here, the positive pair, \mathbf{z}'_i and \mathbf{z}_i , are extracted from the augmented versions of the same sample x_i .

Local self-supervised contrastive loss. Though L_{global}^{ss} (Eq. (1)) favors transferable representations based on global feature vector \mathbf{h}_i , it might ignore some local discriminative information in feature map $\hat{\mathbf{x}}_i$ which could be beneficial in meta-testing. Inspired by [1,3,16,32], we compute self-supervised contrastive loss at the local level. Unlike previous approaches, we leverage map-map and vector-map modules to boost the robustness and generalizability of the representations. The **map-map module** is illustrated in Fig. 3a. Specifically, we use three f_q, f_k, f_v spatial projection heads to project local feature map $\hat{\mathbf{x}}_i$ into the

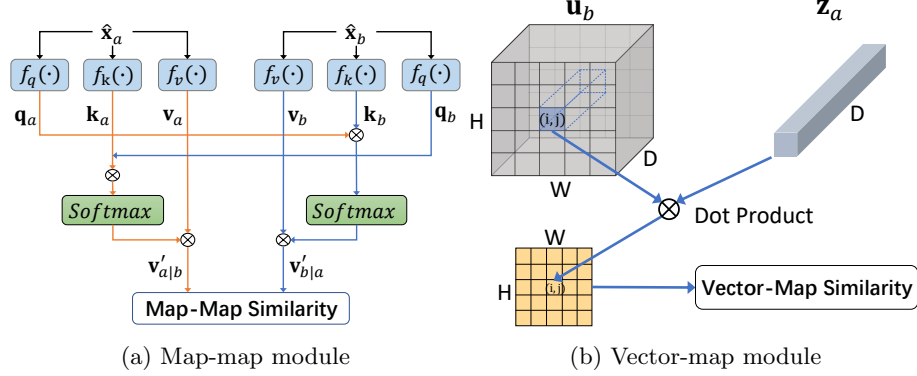


Fig. 3: Map-map and vector-map modules. (a) Operation \otimes denotes dot product. Local feature maps $\hat{\mathbf{x}}_a$ and $\hat{\mathbf{x}}_b$ share three spatial projection heads $f_q(\cdot)$, $f_k(\cdot)$ and $f_v(\cdot)$. We first align $\hat{\mathbf{x}}_a$ with \mathbf{q}_b and then align $\hat{\mathbf{x}}_b$ with \mathbf{q}_a . (b) We obtain \mathbf{u}_b by adding a FC layer after $\hat{\mathbf{x}}_b$ and \mathbf{z}_a denotes the projected vector from $proj(\hat{\mathbf{x}}_a)$.

query $\mathbf{q}_i = f_q(\hat{\mathbf{x}}_i)$, key $\mathbf{k}_i = f_k(\hat{\mathbf{x}}_i)$ and value $\mathbf{v}_i = f_v(\hat{\mathbf{x}}_i)$, respectively, where $\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i \in \mathbb{R}^{HW \times D}$. For a pair of local feature map $\hat{\mathbf{x}}_a$ and $\hat{\mathbf{x}}_b$, we align the $\hat{\mathbf{x}}_a$ with $\hat{\mathbf{x}}_b$ to obtain $\mathbf{v}'_{a|b} = \text{softmax}\left(\frac{\mathbf{q}_b \mathbf{k}_a^\top}{\sqrt{d}}\right) \mathbf{v}_a$, and align $\hat{\mathbf{x}}_b$ with $\hat{\mathbf{x}}_a$ to obtain $\mathbf{v}'_{b|a} = \text{softmax}\left(\frac{\mathbf{q}_a \mathbf{k}_b^\top}{\sqrt{d}}\right) \mathbf{v}_b$. After l_2 normalization on each position (i, j) of the aligned results, we can compute the similarity between the two local feature maps $\hat{\mathbf{x}}_a$ and $\hat{\mathbf{x}}_b$ as follows:

$$\text{sim}_1(\hat{\mathbf{x}}_a, \hat{\mathbf{x}}_b) = \frac{1}{HW} \sum_{1 \leq i \leq H, 1 \leq j \leq W} (\mathbf{v}'_{a|b})_{ij}^\top (\mathbf{v}'_{b|a})_{ij}. \quad (2)$$

Basically, Eq. (2) calculates the summation of the element-wise product of two feature maps $\mathbf{v}'_{a|b}$ and $\mathbf{v}'_{b|a} \in \mathbb{R}^{HW \times D}$. The self-supervised contrastive loss based on pairwise feature maps can be computed as follows:

$$L_{\text{map-map}}^{ss} = - \sum_{i=1}^{2N} \log \frac{\exp(\text{sim}_1(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}'_i) / \tau_2)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} \exp(\text{sim}_1(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}_j) / \tau_2)}, \quad (3)$$

where $(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}'_i)$ is a positive pair and τ_2 denotes a temperature parameter, and $\mathbb{1}$ is an indicator function. Meanwhile, we adopt **vector-map module** to further exploit the local contrastive information between instances, which is shown in Fig. 3b. In specific, we use a fully connected (FC) layer to obtain $\mathbf{u}_i = g(\hat{\mathbf{x}}_i) = \sigma(\mathbf{W}\hat{\mathbf{x}}_i) \in \mathbb{R}^{D \times HW}$, where σ is a ReLU nonlinearity. We can compute the similarity between a contrast pair as $\text{sim}_2(\hat{\mathbf{x}}_a, \hat{\mathbf{x}}_b) = \frac{1}{HW} \sum_{1 \leq i \leq H, 1 \leq j \leq W} (\mathbf{u}_b)_{ij}^\top \mathbf{z}_a$, where \mathbf{z}_a is the projected vector of $\hat{\mathbf{x}}_a$. The self-supervised contrastive loss based on pairs of feature vectors and feature maps can be computed as follows:

$$L_{\text{vec-map}}^{ss} = - \sum_{i=1}^{2N} \log \frac{\exp(\text{sim}_2(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}'_i) / \tau_3)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} \exp(\text{sim}_2(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}_j) / \tau_3)}, \quad (4)$$

where τ_3 and $\mathbb{1}$ act the same as in Eq. (3). Therefore, the local self-supervised contrastive loss can be defined as:

$$L_{local}^{ss} = L_{vec-map}^{ss} + L_{map-map}^{ss}. \quad (5)$$

Global supervised contrastive loss. To exploit the correlations among individual instances and the correlations among different instances from the same category, we also adopt supervised contrastive loss [19] as follows:

$$L_{global}^s = \sum_{i=1}^{2N} \frac{1}{|P(i)|} \sum_{p \in P(i)} L_{ip}, \quad (6)$$

where $L_{ip} = -\log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_p / \tau_4)}{\sum_{j=1}^{2N} \mathbb{1}_{j \neq i} \exp(\mathbf{z}_i \cdot \mathbf{z}_j / \tau_4)}$, and τ_4 is a temperature parameter. The set $P(i)$ contains the indexes of samples with the same label as x_i in the augmented batch except for index i .

To summarize, we minimize the following loss during pre-training:

$$L_{pre} = L_{CE} + \alpha_1 L_{global}^{ss} + \alpha_2 L_{local}^{ss} + \alpha_3 L_{global}^s, \quad (7)$$

where L_{CE} is the CE loss, and α_1, α_2 and α_3 are balance scalars. By optimizing L_{pre} , we promote the discriminability and generalizability of the representations, which is crucial for the following stage.

3.4 Meta-training

Cross-view Episodic Training. In order to capture information about the high-level concept of a shared context, a common practice in contrastive learning [3,14,30,50] is to maximize the mutual information between features extracted from multiple views of the shared context. Intuitively, given an episode $E = \{\mathcal{S}, \mathcal{Q}\}$ in episodic meta-learning, we can obtain two episodes $E_1 = \{\mathcal{S}_1, \mathcal{Q}_1\}$ and $E_2 = \{\mathcal{S}_2, \mathcal{Q}_2\}$ by applying two different data augmentation strategies on E respectively. Here, we consider E a shared context and treat the two augmented episodes as its two views. Inspired by the above idea, we propose a cross-view episodic training mechanism with the aim of forcing representations to learn meta-knowledge that can play a key role across various few-shot classification tasks. Specifically, we use the pre-trained feature extractor f_ϕ in Sect. 3.3 followed by a GAP to each data point x_i in both the episode E_1 and E_2 , and derive the corresponding global vector $\mathbf{h}_i \in \mathbb{R}^C$. We then separately compute a prototype \mathbf{c}_r^k for category k in support sets \mathcal{S}_1 and \mathcal{S}_2 as follows:

$$\mathbf{c}_r^k = \frac{1}{|\mathcal{S}_r^k|} \sum_{(\mathbf{h}_i, y_i) \in \mathcal{S}_r^k} \mathbf{h}_i, \quad (8)$$

where $r \in \{1, 2\}$ and \mathcal{S}_r^k denotes the set of data points belonging to class $k \in \{1, 2, \dots, M\}$ from r -th support set. We denote the task-specific module (multi-head attention [46]) proposed in baseline [54] as $Attn(\cdot)$ and fix the number of heads to 1. With the task-specific module, we obtain an aligned prototype set

$\mathcal{T}(r) = \{(\mathbf{c}_r^k)'\}_{k=1}^M = \text{Attn}(\mathbf{c}_r^1, \mathbf{c}_r^2, \dots, \mathbf{c}_r^M)$ for each support set. Then, based on the aligned prototypes, a probability distribution of a data point x_i in the query set \mathcal{Q}_r over M classes is defined as:

$$P(y = k | \mathbf{h}_i, \mathcal{T}(r)) = \frac{\exp(-d(\mathbf{h}_i, (\mathbf{c}_r^k)'))}{\sum_{j=1}^M \exp(-d(\mathbf{h}_i, (\mathbf{c}_r^j)'))}, \quad (9)$$

where $d(\cdot, \cdot)$ denotes Euclidean distance. Therefore the loss of the nearest centroid classifier on an episode can be computed as:

$$L_{mn} = \frac{1}{|\mathcal{Q}_m|} \sum_{(\mathbf{h}_i, y_i) \in \mathcal{Q}_m} -\log P(y = k | \mathbf{h}_i, \mathcal{T}(n)), \quad (10)$$

where $m, n \in \{1, 2\}$. Obviously different from the original episodic training, we classify \mathcal{Q}_1 on \mathcal{S}_1 and \mathcal{S}_2 respectively, and do the same for \mathcal{Q}_2 . Eq.(10) minimizes the differences between two views of instances of the same category. The whole process is illustrated in Fig. 2. Therefore, we computed the cross-view classification loss as follows:

$$L_{meta} = \frac{1}{4} \sum_{m,n} L_{mn}. \quad (11)$$

Distance-scaled Contrastive Loss. Since contrastive learning approaches work solely at the instance level, it is superficial to simply add contrastive loss into meta-training without taking full advantage of the episodic training mechanism catered for FSL. To better apply contrastive learning to the meta-training, we perform query instance discrimination between two views of the shared episodes. Specifically, inspired by [4], we further inherit the pre-trained projection head $proj(\cdot)$ in Sect. 3.3 to map each sample x_i in the two episodes E_1 and E_2 into a projected vector $\mathbf{z}_i \in \mathbb{R}^D$. We similarly obtain the prototype \mathbf{o}_r^k by averaging the projected vector of the same classes, that is $\mathbf{o}_r^k = \frac{1}{|\mathcal{S}_r^k|} \sum_{(\mathbf{z}_i, y_i) \in \mathcal{S}_r^k} \mathbf{z}_i$ where r is the same as in Eq. (8). For each query vector \mathbf{z}_i in \mathcal{Q}_1 and \mathcal{Q}_2 , we reconstruct its positive sample set by using corresponding augmented version \mathbf{z}'_i and the samples of class y_i in both support sets. Therefore, we reformulate the supervised contrastive loss [19] in the form of episodic training as follows:

$$L(\mathbf{z}_i) = - \sum_{\mathbf{z}_H \in H(\mathbf{z}_i)} \log \frac{\lambda_{\mathbf{z}_i \mathbf{z}_H} \exp(\mathbf{z}_i \cdot \mathbf{z}_H / \tau_5)}{\sum_{\mathbf{z}_A \in A(\mathbf{z}_i)} \lambda_{\mathbf{z}_i \mathbf{z}_A} \exp(\mathbf{z}_i \cdot \mathbf{z}_A / \tau_5)}. \quad (12)$$

Here, operation \cdot means the inner product between features after l_2 normalization. τ_5 is a temperature parameter. The $\lambda_{\mathbf{z}_i \mathbf{z}_j} = 2 - \text{dist}(\mathbf{z}_i, \mathbf{z}_j)$ is the coefficient that reflects the distance relationship between \mathbf{z}_i and \mathbf{z}_j , where $\text{dist}(\cdot, \cdot)$ refers to cosine similarity. $H(\mathbf{z}_i) = \{\mathbf{z}'_i\} \cup \mathcal{S}_1^{y_i} \cup \mathcal{S}_2^{y_i}$ is the positive set of \mathbf{z}_i , and $A(\mathbf{z}_i) = \{\mathbf{z}'_i\} \cup \mathcal{S}_1 \cup \mathcal{S}_2 \cup \{\mathbf{o}_1^k\}_{k=1}^M \cup \{\mathbf{o}_2^k\}_{k=1}^M$. The distance coefficient $\lambda_{\mathbf{z}_i \mathbf{z}_j}$ and additional prototypes \mathbf{o}_r^k on both views of the original episode are introduced to reduce the similarities of queries to their positives. Thus the model can learn more discriminative representations adapted to different tasks. Then we compute

the distance-scaled contrastive loss as:

$$L_{info} = \sum_{\mathbf{z}_i \in \mathcal{Q}_1 \cup \mathcal{Q}_2} \frac{1}{|H(\mathbf{z}_i)|} L(\mathbf{z}_i). \quad (13)$$

By optimizing L_{info} , we force the representations to capture information about instance discrimination episodically and learn the interrelationships among samples of the same category in cross-view episodes.

Objective in Meta-training. In meta-training, we mainly build two losses as in Eq. (11) and Eq. (13) in order to enhance the transferability and discriminative ability of representations. Then the total objective function in meta-training is defined as:

$$L_{total} = L_{meta} + \beta L_{info}, \quad (14)$$

where β is a balance scalar and will be detailed in the following section.

4 Experiments

4.1 Datasets and Setup

Datasets. We evaluate our method on three popular benchmark datasets. The *miniImageNet* dataset [47] contains 100 classes with 600 images per class, and these classes are divided into 64, 16 and 20 for the training, validation and test sets, respectively. The *tieredImageNet* dataset [35] contains 608 classes grouped in 34 high-level categories with 779,165 images, where 351 classes are used for training, 97 for validation and 160 for testing. The *CIFAR-FS* dataset [2] contains 100 classes with 600 images per class. These classes are split into the training, validation and test sets in proportions of 64, 16 and 20.

Implementation Details. Following the baseline [54], we use ResNet-12 as backbone and a multihead attention as the task-specific module (number of heads is 1). All projection heads in our method has the same structure as in [3]. We adopt SGD optimizer with a weight decay of 5e-4 and a momentum of 0.9 for both the pre-training and meta-training stages. During pre-training, the learning rate is initialized to be 0.1 and adapted via cosine learning rate scheduler after warming up. Temperature parameters $\tau_{1,2,3,4}$ are set to 0.1 and the balance scalars $\alpha_{1,2,3}$ are set to 1.0. During meta-training, temperature τ_5 are set to 0.1. We use StepLR with a step size of 40 and gamma of 0.5 and set $\beta = 0.01$ for 1-shot. For 5-shot, StepLR is used with a step size of 50 and gamma of 0.5, and the β is set to 0.1.

Evaluation. We follow the 5-way 1-shot and 5-way 5-shot few-shot classification tasks. In meta-testing, our method (inductive) simply classify the query samples by computing the euclidean distance between the prototypes and query samples. We randomly sample 2000 episodes from test set in meta-testing with 15 query images per class and report the mean accuracy together with corresponding 95% confidence interval.

Data Augmentation. For all datasets, we empirically find that standard [54,55] and SimCLR-style [3,14] data augmentation strategies work best in pre-training

Table 1: The average 5-way few-shot classification accuracies(%) with 95% confidence interval on miniImageNet and tieredImageNet.

Method	Backbone	miniImageNet		tieredImageNet		
		1-shot	5-shot	1-shot	5-shot	
MAML [10]	CONV-4	48.70 ± 1.75	63.11 ± 0.92	—	—	
RelationNets [42]		50.44 ± 0.82	65.32 ± 0.70	54.48 ± 0.93	71.32 ± 0.78	
MatchingNets [47]		48.14 ± 0.78	63.48 ± 0.66	—	—	
ProtoNets [40]		44.42 ± 0.84	64.24 ± 0.72	53.31 ± 0.89	72.69 ± 0.74	
LEO [36]	WRN-28-10	61.76 ± 0.08	77.59 ± 0.12	66.33 ± 0.05	82.06 ± 0.08	
CC+rot [11]		62.93 ± 0.45	79.87 ± 0.33	62.93 ± 0.45	79.87 ± 0.33	
wDAE [12]		61.07 ± 0.15	76.75 ± 0.11	68.18 ± 0.16	83.09 ± 0.12	
PSST [7]		64.16 ± 0.44	80.64 ± 0.32	—	—	
TADAM [31]	ResNet-12	58.5 ± 0.3	76.7 ± 0.3	—	—	
MetaOptNet [22]		62.64 ± 0.61	78.63 ± 0.46	65.99 ± 0.72	81.56 ± 0.53	
DeepEMD [55]		65.91 ± 0.82	82.41 ± 0.56	71.16 ± 0.87	86.03 ± 0.58	
CAN [17]		63.85 ± 0.48	79.44 ± 0.34	69.89 ± 0.51	84.23 ± 0.37	
FEAT [54]		66.78 ± 0.20	82.05 ± 0.14	70.80 ± 0.23	84.79 ± 0.16	
RFS [45]		62.02 ± 0.63	79.64 ± 0.44	69.74 ± 0.72	84.41 ± 0.55	
InfoPatch [25]		67.67 ± 0.45	82.44 ± 0.31	71.51 ± 0.52	85.44 ± 0.35	
DMF [52]		67.76 ± 0.46	82.71 ± 0.31	71.89 ± 0.52	85.96 ± 0.35	
RENet [18]		67.60 ± 0.44	82.58 ± 0.30	71.61 ± 0.51	85.28 ± 0.35	
BML [56]		67.04 ± 0.63	83.63 ± 0.29	68.99 ± 0.50	85.49 ± 0.34	
PAL [26]		69.37 ± 0.64	84.40 ± 0.44	72.25 ± 0.72	86.95 ± 0.47	
TPMN [49]		67.64 ± 0.63	83.44 ± 0.43	72.24 ± 0.70	86.55 ± 0.63	
Ours		ResNet-12	70.19 ± 0.46	84.66 ± 0.29	72.62 ± 0.51	86.62 ± 0.33

and meta-training. During meta-testing, no data augmentation strategy is used. The image transformations used in standard strategy include randomresizedcrop, colorJitter and randomhorizontalflip, while SimCLR-style strategy contains randomresizedcrop, randomhorizontalflip, randomcolorjitter and randomgrayscale.

4.2 Main Results

In this subsection, we compare our method with competitors on three mainstream FSL datasets and the results are reported in Table 1 and Table 2. We can observe that our proposed method consistently achieves competitive results compared to the current state-of-the-art (SOTA) FSL methods on both the 5-way 1-shot and 5-way 5-shot tasks. For the miniImageNet dataset (Table 1), our proposed method outperforms the current best results by 0.82% in the 1-shot task and 0.26% in 5-shot task. For the tieredImageNet dataset (Table 1), our method improves over the current SOTA method by 0.37% for 1-shot and achieves the second best 5-shot result. Note that our method outperforms the original FEAT by 1.82% for 1-shot and 1.83% for 5-shot on tieredImageNet. For the CIFAR-FS dataset (Table 2), our method surpasses the current SOTA by 0.46% and 0.19% in the 1-shot and 5-shot tasks, respectively. Compared to those methods [7,11,25,26], our proposed method works better on most datasets. The consistent and competitive results on the three datasets indicate that our

Table 2: The average 5-way few-shot classification accuracies(%) with 95% confidence interval on CIFAR-FS. * results used our implementation.

Method	Backbone	1-shot	5-shot
Ravichandran <i>et al.</i> [34]	CONV-4	55.14 \pm 0.48	71.66 \pm 0.39
ConstellationNet [53]		69.3 \pm 0.3	82.7 \pm 0.2
CC+rot [11]	WRN-28-10	75.38 \pm 0.31	87.25 \pm 0.21
PSST [7]		77.02 \pm 0.38	88.45 \pm 0.35
Ravichandran <i>et al.</i> [34]	ResNet-12	69.15 \pm -	84.7 \pm -
MetaOptNet [22]		72.0 \pm 0.7	84.2 \pm 0.5
Kim <i>et al.</i> [20]		73.51 \pm 0.92	85.65 \pm 0.65
FEAT* [54]		75.41 \pm 0.21	87.32 \pm 0.15
RFS [45]		73.9 \pm 0.8	86.9 \pm 0.5
ConstellationNet [53]		75.4 \pm 0.2	86.8 \pm 0.2
RENet [18]		74.51 \pm 0.46	86.60 \pm 0.32
BML [56]		73.45 \pm 0.47	88.04 \pm 0.33
PAL [26]		77.1 \pm 0.7	88.0 \pm 0.5
TPMN [49]		75.5 \pm 0.9	87.2 \pm 0.6
Ours		ResNet-12	77.56 \pm 0.47

method can learn more transferable representations by incorporating contrastive learning in both the pre-training and meta-training stages.

4.3 Ablation Study

In this subsection, we study the effectiveness of different components in our method on three datasets. The results in Table 3 show a significant improvement in the performance of our proposed method compared to the baseline pre-trained by CE loss (L_{CE}) only. Specifically, using L_{CE} and all contrastive losses L_{CL} in pre-training improves the accuracy by an average of 1.66% (1-shot) and 1.29% (5-shot) on the three datasets. This allows the representations to generalize better, rather than just focusing on the information only needed for the classification on base classes. We then apply the CVET and L_{info} respectively, and the consistent improvements suggest that CVET and L_{info} are effective. The results jointly based on CVET and L_{info} are further enhanced by 0.95% (1-shot) and 0.72% (5-shot) in average, indicating that our full method increases the transferability of representations on novel classes. Note that the gains obtained by CVET and L_{info} are relatively low on miniImageNet, as the number of training steps on this dataset is very small, leading to a sufficiently narrow gap between two different views of the same image at the end of pre-training. Additionally, more ablation experiments on the effectiveness of all parts in L_{CL} are included in the supplementary material.

4.4 Further Analysis

Data augmentation. We investigate the impact of employing a different data augmentation strategy while maintaining the use of the standard one. Table 4

Table 3: Ablation experiments on miniImageNet. L_{CE} means the baseline model is pre-trained with CE loss. L_{CL} denotes all contrastive losses used in pre-training. CVET and L_{info} are only available in meta-training.

L_{CE}	L_{CL}	CVET	L_{info}	miniImageNet		tieredImageNet		CIFAR-FS	
				1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
✓				66.58 ± 0.46	81.92 ± 0.31	70.41 ± 0.51	84.69 ± 0.36	75.54 ± 0.48	87.28 ± 0.32
✓	✓			69.53 ± 0.47	84.33 ± 0.29	71.83 ± 0.51	85.64 ± 0.35	76.15 ± 0.47	87.79 ± 0.33
✓	✓	✓		69.89 ± 0.46	84.43 ± 0.29	72.39 ± 0.52	86.07 ± 0.35	77.02 ± 0.48	88.28 ± 0.32
✓	✓		✓	69.78 ± 0.46	84.46 ± 0.29	72.51 ± 0.52	86.23 ± 0.34	77.37 ± 0.49	88.48 ± 0.32
✓	✓	✓	✓	70.19 ± 0.46	84.66 ± 0.29	72.62 ± 0.51	86.62 ± 0.33	77.56 ± 0.47	88.64 ± 0.31

Table 4: Comparison experiments of different data augmentation strategies on miniImageNet.

Augmentation	1-shot	5-shot
Standard [54,55]	67.84 ± 0.47	82.82 ± 0.31
SimCLR-style [3,14]	70.19 ± 0.46	84.66 ± 0.29
AutoAugment [8]	68.45 ± 0.46	83.94 ± 0.30
StackedRandAug [44]	68.31 ± 0.46	83.13 ± 0.31

shows that the SimCLR-style [3,14] strategy works best while the AutoAugment [8] and StackedRandAug[44] perform slightly worse. However, all three strategies are better than the standard one. We believe that the standard strategy lacks effective random image transformations and thus always produces two similar views of the same image during the training process. Conversely, the AutoAugment and StackedRandAug strategies that intensely change the original image may lead to excessive bias between two views of the same image, thereby increasing the difficulty of performing contrastive learning.

Visualization. We give visualization to validate the transferability of representations produced by our framework on novel classes. Specifically, we randomly sample 5 classes from the meta-test set of miniImageNet with 100 images per class and obtain embeddings of all images using ProtoNet [40], ProtoNet+Ours, baseline and Ours (baseline + our framework), respectively. Then we use t-SNE [27] to project these embeddings into 2-dimensional space as shown in Fig. 4. The distributions of the embeddings obtained by ProtoNet + Ours and baseline + our framework are more separable and the class boundaries more precise and compact (see Fig. 4b, Fig. 4d vs. Fig. 4a, Fig. 4c). The visualization results indicate that our proposed framework generates more transferable and discriminative representations on novel classes. The complete visualization results are available in the supplementary material.

Method Combination. We combine our proposed framework with the two-stage FSL methods ProtoNet [40] and baseline (FEAT [54]), respectively. The results are shown in Table 5. Integration with our framework improves the accuracy of ProtoNet by 2.74% (1-shot) and 1.39% (5-shot). Similarly, our framework improves the 1-shot accuracy of baseline by 3.96% and its 5-shot accuracy by

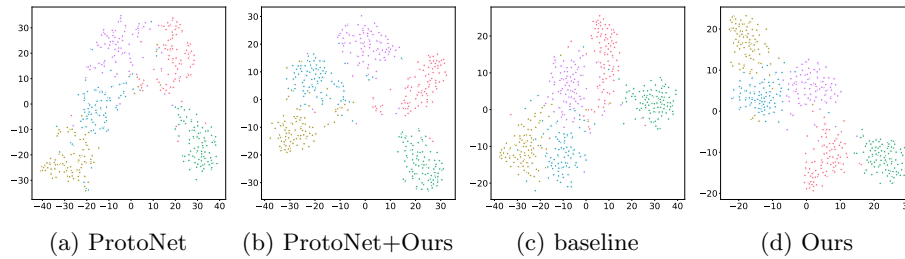


Fig. 4: Visualization of 100 randomly sampled images for each of the 5 meta-test classes from miniImageNet by t-SNE [27].

Table 5: Method combinations on miniImageNet.

Model	1-shot	5-shot
ProtoNet	63.77 ± 0.47	80.58 ± 0.32
ProtoNet+CL	66.17 ± 0.46	81.73 ± 0.30
ProtoNet+Ours	66.51 ± 0.47	81.97 ± 0.30
baseline	66.23 ± 0.47	80.81 ± 0.33
baseline+CL	69.53 ± 0.47	84.33 ± 0.29
Ours	70.19 ± 0.46	84.66 ± 0.29

3.85%. The results in Fig. 4 and Table 5 show that our framework can be applied to other two-stage FSL methods and effectively enhance the performance.

5 Conclusions

In this paper, we apply contrastive learning to the two-stage training paradigm of FSL to alleviate the limitation on generalizability of representations. Concretely, we use vector-map and map-map modules to incorporate self-supervised contrastive losses in the pre-training stage. We further propose a CVET strategy and a distance-scaled contrastive loss to extend contrastive learning to the meta-training stage effectively. The comprehensive experimental results show that our proposed method achieves competitive performance on three well-known FSL datasets miniImageNet, tieredImageNet and CIFAR-FS.

Acknowledgement. This work are supported by: (i) National Natural Science Foundation of China (Grant No. 62072318 and No. 62172285); (ii) Natural Science Foundation of Guangdong Province of China (Grant No. 2021A1515012014); (iii) Science and Technology Planning Project of Shenzhen Municipality (Grant No. JCYJ20190808172007500).

References

1. Bachman, P., Hjelm, R.D., Buchwalter, W.: Learning representations by maximizing mutual information across views. In: NIPS. pp. 15509–15519 (2019)
2. Bertinetto, L., Henriques, J.F., Torr, P.H.S., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. In: ICLR (2019)
3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. In: ICML. pp. 1597–1607 (2020)
4. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.E.: Big self-supervised models are strong semi-supervised learners. In: NIPS. pp. 22243–22255 (2020)
5. Chen, W., Liu, Y., Kira, Z., Wang, Y.F., Huang, J.: A closer look at few-shot classification. In: ICLR (2019)
6. Chen, Y., Liu, Z., Xu, H., Darrell, T., Wang, X.: Meta-baseline: Exploring simple meta-learning for few-shot learning. In: ICCV. pp. 9062–9071 (2021)
7. Chen, Z., Ge, J., Zhan, H., Huang, S., Wang, D.: Pareto self-supervised training for few-shot learning. In: CVPR. pp. 13663–13672 (2021)
8. Cubuk, E.D., Zoph, B., Mané, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: CVPR. pp. 113–123 (2019)
9. Doersch, C., Gupta, A., Zisserman, A.: Crosstransformers: spatially-aware few-shot transfer. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) NIPS. pp. 21981–21993 (2020)
10. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML. pp. 1126–1135 (2017)
11. Gidaris, S., Bursuc, A., Komodakis, N., Pérez, P., Cord, M.: Boosting few-shot visual learning with self-supervision. In: ICCV. pp. 8058–8067 (2019)
12. Gidaris, S., Komodakis, N.: Generating classification weights with GNN denoising autoencoders for few-shot learning. In: CVPR. pp. 21–30 (2019)
13. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: ICLR (2018)
14. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.B.: Momentum contrast for unsupervised visual representation learning. In: CVPR. pp. 9726–9735 (2020)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
16. Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. In: ICLR (2019)
17. Hou, R., Chang, H., Ma, B., Shan, S., Chen, X.: Cross attention network for few-shot classification. In: NIPS. pp. 4005–4016 (2019)
18. Kang, D., Kwon, H., Min, J., Cho, M.: Relational embedding for few-shot classification. In: ICCV. pp. 8822–8833 (2021)
19. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschiot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: NIPS. pp. 18661–18673 (2020)
20. Kim, J., Kim, H., Kim, G.: Model-agnostic boundary-adversarial sampling for test-time generalization in few-shot learning. In: ECCV. pp. 599–617 (2020)
21. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. pp. 1106–1114 (2012)
22. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: CVPR. pp. 10657–10665 (2019)

23. Li, W., Wang, L., Xu, J., Huo, J., Gao, Y., Luo, J.: Revisiting local descriptor based image-to-class measure for few-shot learning. In: CVPR. pp. 7260–7268 (2019)
24. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few shot learning. arXiv preprint arXiv:1707.09835 (2017)
25. Liu, C., Fu, Y., Xu, C., Yang, S., Li, J., Wang, C., Zhang, L.: Learning a few-shot embedding model with contrastive learning. In: AAAI. pp. 8635–8643 (2021)
26. Ma, J., Xie, H., Han, G., Chang, S.F., Galstyan, A., Abd-Elmageed, W.: Partner-assisted learning for few-shot image classification. In: ICCV. pp. 10573–10582 (2021)
27. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *JMLR* **9**(11) (2008)
28. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. arXiv preprint arXiv:1803.02999 (2018)
29. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: ECCV. pp. 69–84 (2016)
30. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
31. Oreshkin, B.N., López, P.R., Lacoste, A.: TADAM: task dependent adaptive metric for improved few-shot learning. In: NIPS. pp. 719–729 (2018)
32. Ouali, Y., Hudelot, C., Tami, M.: Spatial contrastive learning for few-shot classification. In: ECML-PKDD. pp. 671–686 (2021)
33. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2017)
34. Ravichandran, A., Bhotika, R., Soatto, S.: Few-shot learning with embedded class models and shot-free meta training. In: ICCV. pp. 331–339 (2019)
35. Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: ICLR (2018)
36. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: ICLR (2019)
37. Satorras, V.G., Estrach, J.B.: Few-shot learning with graph neural networks. In: ICLR (2018)
38. Shen, Z., Liu, Z., Qin, J., Savvides, M., Cheng, K.: Partial is better than all: Revisiting fine-tuning strategy for few-shot learning. In: AAAI. pp. 9594–9602 (2021)
39. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
40. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. In: NIPS. pp. 4077–4087 (2017)
41. Su, J., Maji, S., Hariharan, B.: When does self-supervision improve few-shot learning? In: ECCV. pp. 645–666 (2020)
42. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H.S., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR. pp. 1199–1208 (2018)
43. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. In: ECCV. pp. 776–794 (2020)
44. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? In: NIPS. pp. 6827–6839 (2020)
45. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J.B., Isola, P.: Rethinking few-shot image classification: A good embedding is all you need? In: ECCV. pp. 266–282 (2020)

46. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NIPS. pp. 5998–6008 (2017)
47. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: NIPS. pp. 3630–3638 (2016)
48. Wang, Y., Chao, W.L., Weinberger, K.Q., van der Maaten, L.: SimpleShot: Revisiting nearest-neighbor classification for few-shot learning. arXiv preprint arXiv:1911.04623 (2019)
49. Wu, J., Zhang, T., Zhang, Y., Wu, F.: Task-aware part mining network for few-shot learning. In: ICCV. pp. 8433–8442 (2021)
50. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR. pp. 3733–3742 (2018)
51. Xing, C., Rostamzadeh, N., Oreshkin, B.N., Pinheiro, P.O.: Adaptive cross-modal few-shot learning. In: NIPS. pp. 4848–4858 (2019)
52. Xu, C., Fu, Y., Liu, C., Wang, C., Li, J., Huang, F., Zhang, L., Xue, X.: Learning dynamic alignment via meta-filter for few-shot learning. In: CVPR. pp. 5182–5191 (2021)
53. Xu, W., Xu, Y., Wang, H., Tu, Z.: Attentional constellation nets for few-shot learning. In: ICLR (2021)
54. Ye, H., Hu, H., Zhan, D., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: CVPR. pp. 8805–8814 (2020)
55. Zhang, C., Cai, Y., Lin, G., Shen, C.: Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In: CVPR. pp. 12200–12210 (2020)
56. Zhou, Z., Qiu, X., Xie, J., Wu, J., Zhang, C.: Binocular mutual learning for improving few-shot classification. In: ICCV. pp. 8402–8411 (2021)