
Contrastive Prototype Learning with Augmented Embeddings for Few-Shot Learning

Yizhao Gao¹

Nanyi Fei²

Guangzhen Liu²

Zhiwu Lu¹

Tao Xiang³

¹Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

²School of Information, Renmin University of China, Beijing, China

³University of Surrey, Guildford, Surrey, United Kingdom

Abstract

Most recent few-shot learning (FSL) methods are based on meta-learning with episodic training. In each meta-training episode, a discriminative feature embedding and/or classifier are first constructed from a support set in an inner loop, and then evaluated in an outer loop using a query set for model updating. This query set sample centered learning objective is however intrinsically limited in addressing the lack of training data problem in the support set. In this paper, a novel contrastive prototype learning with augmented embeddings (CPLAE) model is proposed to overcome this limitation. First, data augmentations are introduced to both the support and query sets with each sample now being represented as an augmented embedding (AE) composed of concatenated embeddings of both the original and augmented versions. Second, a novel support set class prototype centered contrastive loss is proposed for contrastive prototype learning (CPL). With a class prototype as an anchor, CPL aims to pull the query samples of the same class closer and those of different classes further away. This support set sample centered loss is highly complementary to the existing query centered loss, fully exploiting the limited training data in each episode. Extensive experiments on several benchmarks demonstrate that our proposed CPLAE achieves new state-of-the-art.

1 INTRODUCTION

Deep convolutional neural networks (CNNs) [Krizhevsky et al., 2012, He et al., 2016] have witnessed tremendous successes in many visual recognition tasks. However, the powerful learning ability of CNNs depends on a large amount of manually labeled training data. In practice, sufficient

manual annotation is often too costly and may even be infeasible (e.g., for rare object classes). This has severely limited the usefulness of CNNs for real-world applications. Many attempts have been made recently to mitigate such a limitation from the transfer learning perspective, resulting in the popular research line of few-shot learning (FSL) [Li et al., 2003, 2006]. FSL aims to transfer knowledge learned from abundant seen class samples to a set of unseen classes (only with few shots per class).

Most recent FSL methods are based on meta-learning [Vinyals et al., 2016, Snell et al., 2017, Finn et al., 2017, Sung et al., 2018]. That is, they learn an algorithm or model across a set of sampled FSL training/seen tasks, with the objective of making it generalizable to any unseen test tasks. To that end, an episodic training strategy is adopted, i.e., the seen tasks are arranged into learning episodes, each of which contains n classes and k labeled samples per class to simulate the setting for the unseen test tasks. In each episode, the meta-training data is further split into a support set and a query set. Part of the CNN model (e.g., feature embedding subnet, classification layers, or parameter initialization) to be meta-learned is first obtained in an inner loop using the support set. It is then evaluated in an outer loop using the query set for model updating.

These meta-learning based FSL methods differ mainly in which part of the model is meta-learned. Among them, those meta-learning a feature embedding or distance metric have dominated the state-of-the-art. In particular, many of them [Allen et al., 2019, Li et al., 2019b, Afrasiyabi et al., 2020, Ye et al., 2020, Zhang et al., 2020] are based on the prototypical network (ProtoNet) [Snell et al., 2017] for its simplicity and competitive performance with various extensions. Concretely, given a feature embedding network learned from the preceding episode, these ProtoNet-based methods first compute one prototype per class as the support set class mean; these prototypes are then used as a nearest neighbour classifier in the outer loop on the query set to update the feature embedding. In other words, the meta-learning loss is query centred aiming to make sure that each query sample

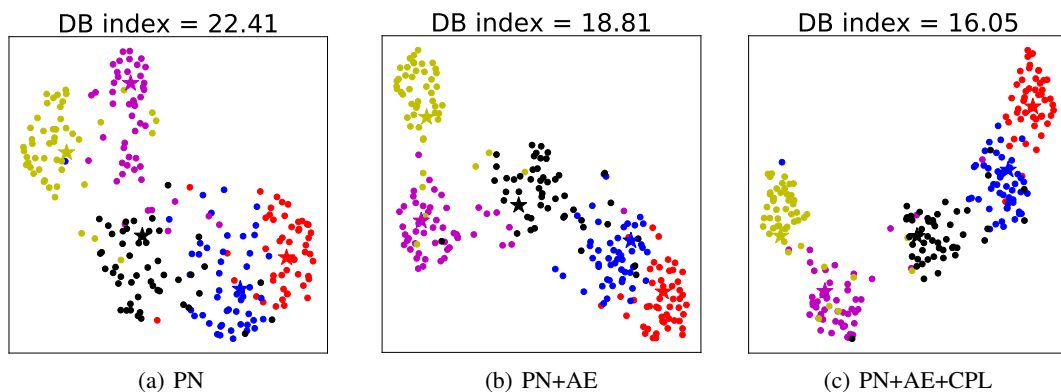


Figure 1: Feature visualization of the same meta-test episode for three FSL models using the UMAP algorithm [McInnes et al., 2018]. PN denotes ProtoNet, which is the prototypical network proposed in [Snell et al., 2017]. Both PN and PN+AE take a query-centered view, while our CPL takes a prototype-centered view. The Davies-Bouldin index (DB index) [Davies and Bouldin, 1979] is used to measure the intra-class variation, which takes a lower value when the data clustering structure is clearer/better.

is close to its corresponding class prototype whilst being further away from other prototypes.

However, this design has severe limitations in addressing a fundamental challenge in FSL, i.e., the lack of support set samples. By definition, each class is only represented by few shots, i.e., n is very small. This problem is actually exacerbated by taking a query-centred only meta-learning loss (with prototype based class representation) that considers the relationship between each query against the k prototypes individually rather than collectively as a distribution.

In this paper, to address the lack of support set sample problem, we propose a novel contrastive prototype learning with augmented embeddings (CPLAE) model for FSL. Our proposed CPLAE has two new components: **(1) Augmented embedding (AE)** – each sample in the support/query set and its three augmented versions are integrated to obtain an augmented embedding. Data augmentation is commonly used in training a CNN for improving its generalization to unseen test data. It has also been considered for FSL [Gidaris et al., 2019, Su et al., 2020, Mangla et al., 2020]. Rather than using augmentation for auxiliary tasks as in existing works, we concatenate the feature embeddings of both the original and augmented versions of each sample to form a richer AE space for meta-learning. **(2) Contrastive prototype learning (CPL)** – Similar to [Allen et al., 2019, Li et al., 2019b, Afrasiyabi et al., 2020, Ye et al., 2020, Zhang et al., 2020], our CPLAE is also based on ProtoNet for meta-learning a feature embedding. Differently, instead of using only query sample centered learning objectives, we additionally introduce a novel support sample centered loss to make full use of the limited training data in each episode. Our CPL loss is a supervised contrastive loss [Khosla et al., 2020] adapted to FSL. More specifically, each prototype is used as an anchor with query set samples of the same class

used as positives and all other query samples as negatives. Contrary and yet complementary to the existing query centered loss that constrains the support set distribution, this support set prototype centered loss regularizes the query set distribution. Combining both losses results in a better embedding space where different classes are more separable (see Figure 1). CPL and AE are integrated seamlessly in our CPLAE in that different AE concatenation orders are applied to the anchor and negatives/positives to further boost the generalization ability of the learned embedding.

Our main contributions are: (1) For the first time, we identify the limitations of existing embedding-based meta-learning methods in dealing with scarce training samples for FSL, caused by adopting only query centered learning objectives. (2) As a remedy, we propose a novel CPLAE model composed of two components (i.e., AE and CPL). Combining AE and supervised contrastive learning seamlessly, our CPL loss enforces a support set centered constraint on the query set sample distribution, thus being complementary to existing query centered losses and effectively making full use of the limited training data. (3) Extensive experiments on several benchmarks demonstrate that our proposed CPLAE achieves new state-of-the-art.

2 RELATED WORK

Few-Shot Learning. Most recent FSL methods follow the meta-learning paradigm. They can be roughly divided into four groups: (1) Embedding/Metric-based methods learn shared task-agnostic embedding spaces/distance metrics or learn task-specific metrics. The former methods either learn an embedding space where a fixed metric (e.g., cosine [Vinyals et al., 2016] or Euclidean distance [Snell et al., 2017]) can be used, or learn a distance metric (e.g., CNN-

based relation modules [Sung et al., 2018, Wu et al., 2019], ridge regression [Bertinetto et al., 2019], and graph neural networks [Satorras and Estrach, 2018, Kim et al., 2019, Yang et al., 2020]). The latter methods learn task-specific metrics [Yoon et al., 2019, Li et al., 2019a, Qiao et al., 2019, Ye et al., 2020, Simon et al., 2020] which can adapt to each unseen new task. (2) Optimization-based methods [Ravi and Larochelle, 2017, Munkhdalai and Yu, 2017, Finn et al., 2017, Nichol et al., 2018, Rusu et al., 2019, Lee et al., 2019] aim to meta-learn an optimizer. Specifically, MAML [Finn et al., 2017] was proposed to learn a good model initialization with seen class data and then quickly adapt it on novel class tasks. Reptile [Nichol et al., 2018] further simplified MAML, and MetaOptNet [Lee et al., 2019] enhanced MAML by replacing the linear classifier with an SVM. (3) Hallucination-based methods [Hariharan and Girshick, 2017, Zhang et al., 2019, Li et al., 2020] aim to learn generators from seen class samples, which are then applied during meta-testing by hallucinating new samples/features using the few shots from unseen classes. (4) Prediction-based methods [Qi et al., 2018, Qiao et al., 2018, Gidaris and Komodakis, 2019, Guo and Cheung, 2020] directly learn to utilize a few labeled samples to predict the parameters of neural networks for few-shot classification.

The state-of-the-art FSL results are mostly achieved by methods from the first group [Allen et al., 2019, Li et al., 2019b, Afrasiyabi et al., 2020, Ye et al., 2020, Zhang et al., 2020], especially those based on ProtoNet [Snell et al., 2017]. Our CPLAE is also an embedding-based method based on ProtoNet. However, armed with augmented embedding (AE) and additionally introducing a support set prototype centered loss, our model is more capable of dealing with the limited training data in FSL, resulting in superior performance (see Sec. 4).

Data Augmentation for FSL. Several recent works [Hsu et al., 2019, Khodadadeh et al., 2019, Antoniou and Storkey, 2019, Qin et al., 2020, Gidaris et al., 2019, Su et al., 2020, Mangla et al., 2020] have utilized data augmentation for meta-learning based FSL. [Hsu et al., 2019, Khodadadeh et al., 2019, Antoniou and Storkey, 2019, Qin et al., 2020] focus on unsupervised FSL, where augmented data samples and their original version are used to form pseudo classes to enable supervised episodic training. For supervised FSL, [Gidaris et al., 2019, Su et al., 2020, Mangla et al., 2020] take a multi-task learning framework where augmented data are used for auxiliary self-supervised pretext tasks (e.g., predicting the rotation angle). Our CPLAE is also a supervised FSL model, but the way data augmentation is used is very different from that in [Gidaris et al., 2019, Su et al., 2020, Mangla et al., 2020]. Specifically, for each sample, we conduct three kinds of image deformations and then input the four images (together with the original one) into a feature embedding network to obtain a concatenated augmented embedding (AE) space with higher dimensionality than the

original embedding space. Different orders of concatenation are further used to formulate our contrastive prototype learning (CPL) loss/objective to boost the generalization ability of the learned embedding.

Contrastive Learning. Contrastive learning (CL) has recently achieved great success in self-supervised learning [van den Oord et al., 2018, Tian et al., 2019, Chen et al., 2020, He et al., 2020] where augmented data creates pseudo classes so that supervised learning can be applied. This has been recently extended to supervised CL [Khosla et al., 2020] where given an instance as anchor, all other instances (original and augmented) of the same classes are positives and the rest as negatives. Our CPL loss is essentially also a supervised CL loss. However, there are vital differences: our anchors are prototypes from the support set and critically CL is seamlessly combined with the proposed AE with different embedding concatenation orders applied to the anchor and positives/negatives respectively to challenge the generalization ability of the learned embedding. Note that ProtoTransfer [Medina et al., 2020] also exploits CL for FSL, but under the unsupervised setting only, rather than our supervised FSL problem.

3 METHODOLOGY

3.1 PROBLEM DEFINITION

Let \mathcal{C}_s denote a set of seen classes and \mathcal{C}_u a set of unseen classes, where $\mathcal{C}_s \cap \mathcal{C}_u = \emptyset$. We are given a large sample set $\mathcal{D}_s = \{(x_i, y_i) | y_i \in \mathcal{C}_s, i = 1, \dots, N_s\}$ from \mathcal{C}_s , and a few-shot sample set $\mathcal{D}_u = \{(x_i, y_i) | y_i \in \mathcal{C}_u, i = 1, \dots, N_u\}$ from \mathcal{C}_u , where x_i is the i -th image in \mathcal{D}_s (or \mathcal{D}_u), y_i is the class label of x_i , and N_s (or N_u) is the number of images in \mathcal{D}_s (or \mathcal{D}_u). Particularly, for the k -shot sample set \mathcal{D}_u , $N_u = k|\mathcal{C}_u|$ (i.e., each class has k labeled images). A test set \mathcal{D}_t from \mathcal{C}_u is also given, where $\mathcal{D}_u \cap \mathcal{D}_t = \emptyset$. The goal of few-shot learning (FSL) is to predict the labels of test images in \mathcal{D}_t by exploiting \mathcal{D}_s and \mathcal{D}_u for training.

3.2 FSL WITH AUGMENTED EMBEDDINGS

Most FSL methods [Finn et al., 2017, Snell et al., 2017, Satorras and Estrach, 2018, Sung et al., 2018, Lee et al., 2019, Kim et al., 2019, Ye et al., 2020] adopt episodic training on the set of seen class samples \mathcal{D}_s and evaluate their models over few-shot classification tasks (i.e., episodes) sampled from the unseen classes \mathcal{C}_u . To form an n -way k -shot episode $e = (\mathcal{S}, \mathcal{Q})$, we first randomly sample a set of n classes \mathcal{C} from \mathcal{C}_s (or \mathcal{C}_u), and then generate a support set $\mathcal{S} = \{(x_i, y_i) | y_i \in \mathcal{C}, i = 1, \dots, n \times k\}$ and a query set $\mathcal{Q} = \{(x_i, y_i) | y_i \in \mathcal{C}, i = 1, \dots, n \times q\}$ ($\mathcal{S} \cap \mathcal{Q} = \emptyset$) by sampling k support and q query samples from each class in \mathcal{C} , respectively.

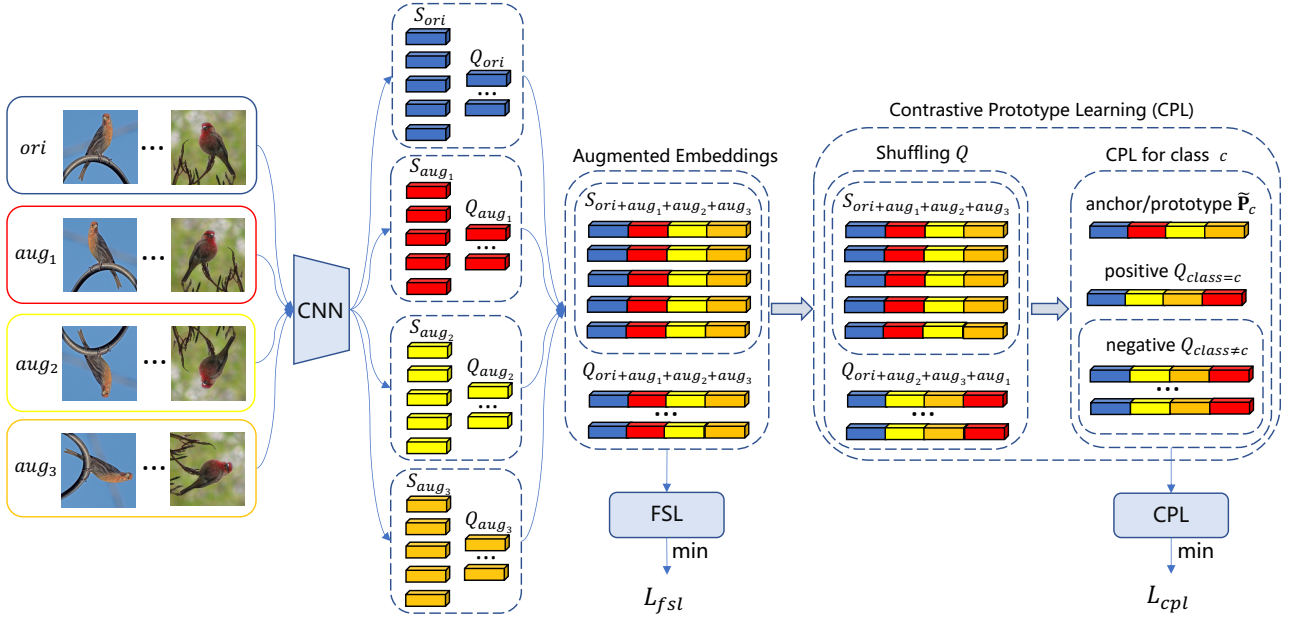


Figure 2: Illustration of our proposed CPLAE model. For each original episode, we conduct three data augmentation methods to generate its three extended episodes. Concretely, samples/embeddings with the subscript ori denote the original ones, while samples/embeddings with the subscripts aug_1 , aug_2 , and aug_3 are obtained by *Horizontal Flip*, *Vertical Flip*, and *Rotation 270°*, respectively. With the augmented embeddings by sample-wise integration, we devise two supervised losses: (1) A standard FSL loss is defined over the augmented embeddings. (2) By shuffling the concatenation order of augmented queries, a novel CPL loss is defined with prototypes as anchors.

We adopt Prototypical Network (ProtoNet) [Snell et al., 2017] as our baseline, which has a feature embedding network and a non-parametric nearest-neighbor classifier. ProtoNet thus only meta-learns the parameters of the embedding network. In each episode, it computes the mean feature embedding of support samples for each class $c \in \mathcal{C}$ as the prototype \mathbf{p}_c :

$$\mathbf{p}_c = \frac{1}{k} \sum_{(x_i, y_i) \in \mathcal{S}} f_\phi(x_i) \cdot I(y_i = c), \quad (1)$$

where f_ϕ denotes the embedding network parameterized by ϕ with an output dimension D , and I denotes the indicator function with its output being 1 if the input is true or 0 otherwise. Once the class prototypes are obtained from the support set, the distance of each query set sample to these prototypes are computed to construct a query centered cross-entropy loss for meta-learning f_ϕ .

To deal with the lack of training data in each episode, we first apply three data augmentation $g^{(1)}$, $g^{(2)}$, and $g^{(3)}$ (e.g., horizontal flip, vertical flip, and rotations) to each image x_i in $\mathcal{S} \cup \mathcal{Q}$ and obtain the corresponding feature embeddings $f_\phi^{(j)}(x_i) = f_\phi(g^{(j)}(x_i))$ ($j = 1, 2, 3$). Together with the feature embedding of the original image, our augmented embedding can be obtained by concatenating the four embeddings (see Figure 2):

$$\tilde{f}_\phi(x_i) = A \left(f_\phi(x_i), f_\phi^{(1)}(x_i), f_\phi^{(2)}(x_i), f_\phi^{(3)}(x_i) \right), \quad (2)$$

where A is an integration function. In this work, we employ the self-attention mechanism [Lin et al., 2017, Vaswani et al., 2017] to update the four input vectors, followed by channel-wise concatenation in order, resulting in a $4D$ -dimensional augmented embedding. We thus have $\tilde{f}_\phi(x_i) \in \mathbb{R}^{4D}$. In this AE space, our prototype for each class $c \in \mathcal{C}$ is obtained as:

$$\tilde{\mathbf{p}}_c = \frac{1}{k} \sum_{(x_i, y_i) \in \mathcal{S}} \tilde{f}_\phi(x_i) \cdot I(y_i = c). \quad (3)$$

For each query sample, by computing the distances to the prototypes, we can formulate the few-shot classification loss over each episode as:

$$L_{fsl} = \frac{1}{q} \sum_{(x_i, y_i) \in \mathcal{Q}} -\log \frac{\exp(-d(\tilde{f}_\phi(x_i), \tilde{\mathbf{p}}_{y_i}))}{\sum_{c \in \mathcal{C}} \exp(-d(\tilde{f}_\phi(x_i), \tilde{\mathbf{p}}_c))}, \quad (4)$$

where $d(\cdot, \cdot)$ denotes the Euclidean distance between two embeddings. Note that this is a conventional FSL loss formulated from a query-centered view. The distribution of the full query set with respect to each prototype is not exploited to regularize the learned feature embedding. This can be achieved by a contrastive prototype learning (CPL) loss formulated from a support-set prototype-centered view.

3.3 CONTRASTIVE PROTOTYPE LEARNING

Our CPL loss is a supervised contrastive learning loss. Different from the conventional un-/self-supervised CL loss, our CPL utilizes the class labels of samples in each episode to construct a few-shot supervised contrastive learning model with augmented embeddings. As illustrated in Figure 2, our main idea of CPL is that: for each class $c \in \mathcal{C}$, we take the prototype $\tilde{\mathbf{p}}_c$ as the anchor, with queries from class c being positive examples and queries from the other classes being negative examples.

Following the common practice in CL where positive examples contain augmented versions of the same training instance/sample, we also adopt augmentation to enrich the training data. However, there is a vital difference: we use the same set of augmentations but vary their concatenation order to produce more nuanced perturbations in the AE space. Concretely, we shuffle the order of the three augmentations and obtain a shuffled augmented embedding $\hat{f}_\phi(x_i)$ of each query sample $x_i \in \mathcal{Q}$ for CPL:

$$\hat{f}_\phi(x_i) = A \left(f_\phi(x_i), f_\phi^{(2)}(x_i), f_\phi^{(3)}(x_i), f_\phi^{(1)}(x_i) \right). \quad (5)$$

Note that the shuffling is only applied to the query samples, not to the support samples/anchors. Also note that in the shuffled concatenation, the original image’s embedding $f_\phi(x_i)$ remain in the first place of $\hat{f}_\phi(x_i)$. We found empirically that once that is fixed, how exactly the other three embeddings are shuffled makes little difference (see the supplementary material for more details).

For each class $c \in \mathcal{C}$, let $\mathcal{P}^{(c)} = \{(x_i, y_i) \in \mathcal{Q} | y_i = c, i = 1, \dots, q\}$ denote the set of positive examples. We then compute the similarity between the anchor/prototype $\tilde{\mathbf{p}}_c$ and each x_i in $\mathcal{P}^{(c)}$ as follows:

$$sim_{c,i}^{(pos)} = \exp(\cos(\tilde{\mathbf{p}}_c, h(\hat{f}_\phi(x_i))))/T, \quad (6)$$

where $h(\cdot)$ is a small neural network projection head that maps representations/embeddings to the space where the contrastive loss is applied (as in [Chen et al., 2020]), $\cos(\cdot, \cdot)$ computes the cosine similarity between two vectors, and T is the temperature parameter. For each positive example $(x_i, y_i) \in \mathcal{P}^{(c)}$, we first randomly sample m ($m \leq q$) query samples from each of the other classes to form the set of negative examples $\mathcal{N}_i^{(c)} = \{(x_t, y_t) \in \mathcal{Q} | y_t \neq c, t = 1, \dots, m(n-1)\}$. We then obtain the similarities for all negative examples:

$$sim_{c,i}^{(neg)} = \sum_{(x_t, y_t) \in \mathcal{N}_i^{(c)}} \exp(\cos(\tilde{\mathbf{p}}_c, h(\hat{f}_\phi(x_t))))/T. \quad (7)$$

The contrastive loss used for CPL is finally given by:

$$L_{cpl} = \frac{1}{nq} \sum_{c \in \mathcal{C}} \sum_{(x_i, y_i) \in \mathcal{P}^{(c)}} -\log \frac{sim_{c,i}^{(pos)}}{sim_{c,i}^{(pos)} + sim_{c,i}^{(neg)}}. \quad (8)$$

Algorithm 1 CPLAE for FSL

Input: Our CPLAE model M_Θ (Θ is the set of parameters)
The seen class training set \mathcal{D}_s
The hyper-parameters λ, T, m

Output: The learned M_Θ^*

- 1: **for all** iteration = 1, 2, \dots , MaxIteration **do**
- 2: Sample an n -way k -shot episode e from \mathcal{D}_s ;
- 3: Obtain $\tilde{f}_\phi(x)$ as the augmented embedding for each sample x from e with Eq. (2);
- 4: Compute L_{fsl} with Eq. (4);
- 5: Obtain $\hat{f}_\phi(x)$ as the shuffled augmented embedding for each query sample x from \mathcal{Q} with Eq. (5);
- 6: Compute L_{cpl} with Eq. (8);
- 7: Compute the total loss L_{total} with Eq. (9);
- 8: Compute the gradients $\nabla_{M_\Theta} L_{total}$;
- 9: Update M_Θ using stochastic gradient descent;
- 10: **end for**
- 11: **return** the found best M_Θ^* .

From this formulation, it is clear that compared to the popular unsupervised CL [Chen et al., 2020], our CPL loss is supervised in that it utilizes the class labels of samples. Compared with existing supervised contrastive losses such as triplet loss [Schroff et al., 2015], its improved version N-pair loss [Sohn, 2016], and the more recent supervised CL loss [Khosla et al., 2020], our CPL loss has two main differences: (1) Our CPL is designed for FSL which takes class prototypes as anchors, while [Schroff et al., 2015, Sohn, 2016, Khosla et al., 2020] take samples as anchors. (2) The contrastive learning is conducted in an AE space with perturbations on the concatenation orders of the augmented feature embeddings to boost the generalization ability of the learned feature embedding.

3.4 LEARNING OBJECTIVES FOR CPLAE

In each training iteration, we randomly sample one n -way k -shot q -query episode $e = (\mathcal{S}, \mathcal{Q})$. For each instance/sample in e , we apply three different data augmentation methods on it, and then integrate the obtained four feature embeddings into one augmented embedding. The few-shot classification loss L_{fsl} is computed with the augmented embeddings according to Eq. (4). Moreover, for all query samples, we shuffle the integrating order of their augmented embeddings and compute the CPL loss L_{cpl} in Eq. (8). The total learning objective for the proposed Contrastive Prototype Learning with Augmented Embeddings (CPLAE) model is finally stated as follows:

$$L_{total} = L_{fsl} + \lambda L_{cpl}, \quad (9)$$

where λ is used to balance the importance of the FSL and CPL losses. In this work, λ is empirically set to 0.1. Our full CPLAE algorithm is outlined in Algorithm 1. Once learned,

Table 1: Comparative results of standard FSL on the two benchmark datasets. The average 5-way few-shot classification accuracies (% , top-1) along with the 95% confidence intervals are reported.

Method	Backbone	<i>miniImageNet</i>		<i>tieredImageNet</i>	
		5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
MatchingNet [Vinyals et al., 2016]	Conv4-64	43.56 ± 0.84	55.31 ± 0.73	–	–
ProtoNet [†] [Snell et al., 2017]	Conv4-64	52.79 ± 0.45	71.23 ± 0.36	53.82 ± 0.48	71.77 ± 0.41
MAM [Finn et al., 2017]	Conv4-64	48.70 ± 1.84	63.10 ± 0.92	51.67 ± 1.81	70.30 ± 0.08
RelationNet [Sung et al., 2018]	Conv4-64	50.40 ± 0.80	65.30 ± 0.70	54.48 ± 0.93	71.32 ± 0.78
IMP [Allen et al., 2019]	Conv4-64	49.60 ± 0.80	68.10 ± 0.80	–	–
DN4 [Li et al., 2019b]	Conv4-64	51.24 ± 0.74	71.02 ± 0.64	–	–
DN PARN [Wu et al., 2019]	Conv4-64	55.22 ± 0.84	71.55 ± 0.66	–	–
PN+rot [Gidaris et al., 2019]	Conv4-64	53.63 ± 0.43	71.70 ± 0.36	–	–
CC+rot [Gidaris et al., 2019]	Conv4-64	54.83 ± 0.43	71.86 ± 0.33	–	–
Centroid [Afrasiyabi et al., 2020]	Conv4-64	53.14 ± 1.06	71.45 ± 0.72	–	–
Neg-Cosine [Liu et al., 2020]	Conv4-64	52.84 ± 0.76	70.41 ± 0.66	–	–
FEAT [Ye et al., 2020]	Conv4-64	55.15 ± 0.20	71.61 ± 0.16	–	–
CPLAE (ours)	Conv4-64	56.83 ± 0.44	74.31 ± 0.34	58.23 ± 0.49	75.12 ± 0.40
ProtoNet [†] [Snell et al., 2017]	Conv4-512	53.52 ± 0.43	73.34 ± 0.36	55.52 ± 0.48	74.07 ± 0.40
MAML [Finn et al., 2017]	Conv4-512	49.33 ± 0.60	65.17 ± 0.49	52.84 ± 0.56	70.91 ± 0.46
Relation Net [Sung et al., 2018]	Conv4-512	50.86 ± 0.57	67.32 ± 0.44	54.69 ± 0.59	72.71 ± 0.43
PN+rot [Gidaris et al., 2019]	Conv4-512	56.02 ± 0.46	74.00 ± 0.35	–	–
CC+rot [Gidaris et al., 2019]	Conv4-512	56.27 ± 0.43	74.30 ± 0.33	–	–
CPLAE (ours)	Conv4-512	57.46 ± 0.43	75.69 ± 0.33	61.56 ± 0.50	80.03 ± 0.38
ProtoNet [†] [Snell et al., 2017]	ResNet-12	62.41 ± 0.44	80.49 ± 0.29	69.63 ± 0.53	84.82 ± 0.36
TADAM [Oreshkin et al., 2018]	ResNet-12	58.50 ± 0.30	76.70 ± 0.38	–	–
MetaOptNet [Lee et al., 2019]	ResNet-12	62.64 ± 0.61	78.63 ± 0.46	65.99 ± 0.72	81.56 ± 0.63
MTL [Sun et al., 2019]	ResNet-12	61.20 ± 1.80	75.50 ± 0.80	65.62 ± 1.80	80.61 ± 0.90
AM3 [Xing et al., 2019]	ResNet-12	65.21 ± 0.49	75.20 ± 0.36	67.23 ± 0.34	78.95 ± 0.22
Shot-Free [Ravichandran et al., 2019]	ResNet-12	59.04 ± 0.43	77.64 ± 0.39	66.87 ± 0.43	82.64 ± 0.43
Neg-Cosine [Liu et al., 2020]	ResNet-12	63.85 ± 0.81	81.57 ± 0.56	–	–
Distill [Tian et al., 2020]	ResNet-12	64.82 ± 0.60	82.14 ± 0.43	71.52 ± 0.69	86.03 ± 0.49
DSN-MR [Simon et al., 2020]	ResNet-12	64.60 ± 0.72	79.51 ± 0.50	67.39 ± 0.82	82.85 ± 0.56
DeepEMD [Zhang et al., 2020]	ResNet-12	65.91 ± 0.82	82.41 ± 0.56	71.16 ± 0.87	86.03 ± 0.58
FEAT [Ye et al., 2020]	ResNet-12	66.78 ± 0.20	82.05 ± 0.14	70.80 ± 0.23	84.79 ± 0.16
CPLAE (ours)	ResNet-12	67.46 ± 0.44	83.22 ± 0.29	72.23 ± 0.50	87.35 ± 0.34

with the optimal model found by our CPLAE algorithm, we randomly sample multiple n -way k -shot meta-test episodes from \mathcal{C}_u for performance evaluation.

4 EXPERIMENTS

4.1 DATASETS AND SETTINGS

Datasets. We select three widely-used benchmarks for evaluation: *miniImageNet* [Vinyals et al., 2016], *tieredImageNet* [Ren et al., 2018], and CUB-200-2011 [Wah et al., 2011]. The *miniImageNet* dataset contains 100 classes from ILSVRC-12 [Russakovsky et al., 2015], with each class having 600 images. We split it into 64 training classes, 16 validation classes, and 20 test classes, as in [Ravi and Larochelle, 2017]. The *tieredImageNet* dataset is a larger subset of ILSVRC-12, which consists of 608 classes and 779,165 images in total. We split it into 351 training classes,

97 validation classes, and 160 test classes, as in [Ren et al., 2018]. Different from the aforementioned two, CUB-200-2011 is a fine-grained classification dataset consisting of 11,778 images from 200 different bird classes. The 200 classes are divided into 100, 50, 50 classes for training, validation and testing, respectively. All images of the datasets are resized to 84×84 before being inputted into the feature embedding networks (i.e., CNNs).

Evaluation Protocols. We make evaluation under 5-way 5-shot/1-shot as in previous works. Each episode has 5 randomly sampled classes from the test split, each of which contains 5 shots (or 1 shot) and 15 queries. We thus have $n = 5$, $k = 5$ or 1, and $q = 15$. Note that since data augmentations can be performed easily (in a fully unsupervised way), we also adopt the augmented embeddings for all images during evaluation, and keep the integration order as in Eq. (2) (i.e., no shuffling is involved). We report average 5-way classification accuracy (% , top-1) over 2,000 meta-test

episodes along with the 95% confidence interval.

Feature Embedding Networks. We adopt three backbones as the feature extractors f_ϕ : Conv4-64 [Vinyals et al., 2016], Conv4-512, and ResNet-12 [He et al., 2016]. They all take the same input image size of 84×84 . Particularly, both Conv4-64 and Conv4-512 consist of 4 convolutional layers: the first three layers are exactly the same but the last layer has different numbers of out channels in the two backbones. Since we use an average pooling layer after the last convolutional layer for each backbone, the output feature dimensions of Conv4-64, Conv4-512, and ResNet-12 are 64, 512, and 640, respectively. We pre-train all three backbones on the training split of each dataset to accelerate the training process, as in [Zhang et al., 2020, Ye et al., 2020, Simon et al., 2020]. With the pre-trained backbones, our CPLAE is then applied in the meta-training stage. For ResNet-12, the stochastic gradient descent (SGD) optimizer is employed with the initial learning rate of $1e-4$, the weight decay of $5e-4$, and the Nesterov momentum of 0.9. For Conv4-64 and Conv4-512, the Adam optimizer [Kingma and Ba, 2015] is adopted with the initial learning rate of $1e-4$.

Implementation Details. In all experiments, our CPLAE is trained for 100 epochs with 100 episodes per epoch, and the learning rate is halved every 20 epochs. The hyper-parameters are selected according to the validation performance of our algorithm. Particularly, for each class $c \in \mathcal{C}$, we sample 6 negative examples (i.e., $m = 6$) from each class in $\mathcal{C} \setminus \{c\}$ for every positive example. While computing the similarity between the anchor and each positive example in Eq. (6), the temperature T is set to 1.

4.2 MAIN RESULTS

For comprehensive comparison, we select a variety of latest/state-of-the-art FSL methods [Ye et al., 2020, Zhang et al., 2020, Simon et al., 2020, Tian et al., 2020, Liu et al., 2020, Afrasiyabi et al., 2020] as well as the strongest SSL+FSL method CC+rot [Gidaris et al., 2019] as the competitors, in addition to the classic/representative baselines (e.g., ProtoNet and MAML). The comparative results of standard/conventional FSL on the *miniImageNet* [Vinyals et al., 2016] and *tieredImageNet* datasets are provided in Table 1. Note that we re-implement our main baseline (i.e., ProtoNet [Snell et al., 2017], denoted with \dagger) with the same hyper-parameters during meta-training for fair comparison. We can see that: (1) With the same backbone (out of the three ones), our CPLAE achieves new state-of-the-art on all datasets under both 1-shot and 5-shot settings, validating the effectiveness of CPL with augmented embeddings (AE). This suggests that our CPLAE has the strongest generalization ability thanks to the introduction of AE and the use of prototype centered contrastive learning. (2) Impressively, our CPLAE with Conv4-64 even outperforms the state-of-the-art competitors with Conv4-512 in all cases. Since the

Table 2: Comparative results for fine-grained FSL on CUB and cross-domain FSL on *miniImageNet* \rightarrow CUB. The training/validation/test of CUB is the same as that used by FEAT [Ye et al., 2020].

Method	5-way 1-shot	5-way 5-shot
CUB:		
MatchingNet	61.16 ± 0.89	72.86 ± 0.70
ProtoNet \dagger	63.72 ± 0.22	81.50 ± 0.15
MAML	55.92 ± 0.95	72.09 ± 0.76
RelationNet	62.45 ± 0.98	76.11 ± 0.69
FEAT	68.87 ± 0.22	82.90 ± 0.15
CPLAE (ours)	69.77 ± 0.50	84.57 ± 0.33
<i>miniImageNet</i> \rightarrow CUB:		
MatchingNet	42.62 ± 0.55	56.53 ± 0.44
ProtoNet \dagger	50.51 ± 0.56	69.28 ± 0.40
MAML	43.59 ± 0.54	54.18 ± 0.41
RelationNet	49.84 ± 0.54	68.98 ± 0.42
FEAT	51.52 ± 0.54	70.16 ± 0.40
CPLAE (ours)	51.67 ± 0.45	71.59 ± 0.40

performance achieved by our CPLAE even surpasses that of the strongest competitor CC+rot [Gidaris et al., 2019] that also utilised data augmentation but with a stronger backbone, our results validate our novel way of using augmentation (self-attention + concatenation) and highlight the importance of support-centered meta-learning loss. (3) The improvements obtained by our CPLAE over the baseline ProtoNet \dagger range from 2.3% to 6.0%, providing direct evidence that both the proposed augmented embeddings and contrastive prototype learning bring significant benefits to FSL (further evidence is provided in ablation study shortly).

We further conduct experiments on CUB-200-2011 (CUB) [Wah et al., 2011] and *miniImageNet* \rightarrow CUB to evaluate our CPLAE model under the fine-grained FSL and cross-domain FSL settings, respectively. CUB is a fine-grained dataset of birds, which has 200 classes and 11,788 images in total. We follow [Ye et al., 2020] and split CUB into 100 training classes, 50 validation classes, and 50 test classes. For cross-domain FSL on *miniImageNet* \rightarrow CUB, the 100 training classes are from *miniImageNet* while the 50 validation and 50 test classes (using the aforementioned split for fine-grained FSL) are from CUB. On both datasets, we use Conv4-64 as the feature extractor. The comparative results under 5-way 1-shot/5-shot settings are shown in Table 2. We can see that our CPLAE model achieves the best results, validating the effectiveness of CPLAE under both fine-grained and cross-domain FSL settings.

4.3 FURTHER EVALUATION

Ablation Study. Our full CPLAE model is trained with two losses: the FSL loss L_{fsl} and the CPL loss L_{cpl} (see Eq. (9)). For L_{fsl} , we adopt an augmented embedding for each sam-

Table 3: Ablation study results for our full CPLAE model (including AE and CPL) on the *miniImageNet* dataset. Conv4-64 is used as the feature extractor.

Method	5-way 1-shot	5-way 5-shot
ProtoNet [†]	52.79 ± 0.45	71.23 ± 0.36
ProtoNet [†] +AE	55.89 ± 0.43	73.43 ± 0.35
ProtoNet [†] +AE+CPL (no shuffling)	56.04 ± 0.44	73.75 ± 0.35
ProtoNet [†] +AE+CPL	56.83 ± 0.44	74.31 ± 0.34

Table 4: Comparison among different choices of the data augmentation methods on the *miniImageNet* dataset. Only the 5-way 5-shot setting is considered, and Conv4-64 is used as the feature extractor.

Data Augmentation Methods					
Horizontal Flip	Vertical Flip	Rotation 90°	Rotation 180°	Rotation 270°	5-way 5-shot
×	×	×	×	×	71.23 ± 0.36
✓	✓				72.96 ± 0.35
✓	✓	✓			73.41 ± 0.35
✓	✓		✓		73.34 ± 0.35
✓	✓			✓	73.43 ± 0.35
✓		✓	✓		73.40 ± 0.34
✓		✓		✓	73.34 ± 0.34
✓			✓	✓	73.28 ± 0.35
	✓	✓	✓		72.62 ± 0.35
	✓	✓		✓	73.14 ± 0.35
	✓		✓	✓	72.90 ± 0.35
		✓	✓	✓	73.14 ± 0.35
✓	✓	✓		✓	72.77 ± 0.36
✓	✓	✓	✓		73.10 ± 0.35

ple, which is obtained by integrating four feature vectors (one from the original image and three from its augmented ones). For L_{cpl} , we devise a novel supervised contrastive loss with the shuffling operation of query samples. To demonstrate the contribution of each main component, we conduct ablative experiments on *miniImageNet* in Table 3, where Conv4-64 is adopted as the backbone. Four methods are compared: (1) ProtoNet[†]: our re-implementation of ProtoNet [Snell et al., 2017]. (2) ProtoNet[†]+AE: ProtoNet trained with augmented embeddings (i.e., trained with only L_{fsl} in Eq. (9)). (3) CPLAE (no-shuffling): Our CPLAE model trained with the total loss in Eq. (9) but without the shuffling operation. (4) CPLAE: our full CPLAE model. The ablation study results in Table 3 show that the augmented embeddings lead to 2–3% improvements (see ProtoNet[†]+AE vs. ProtoNet[†]), and our proposed CPL further improves the performance by about 1% (see CPLAE vs. ProtoNet[†]+AE). In addition, the comparison CPLAE vs. CPLAE (no shuffling) demonstrates the importance of the shuffling operation for our CPL.

Alternative Augmentation Strategies. In Table 4, we compare different choices of the data augmentation methods for our augmented embeddings. We select five common image

Table 5: Comparison to different SSL losses on the *miniImageNet* dataset (with Conv4-64 being the backbone). Notations: PT – SSL based on pretext tasks; CL – SSL via contrastive learning.

Method	5-way 1-shot	5-way 5-shot
ProtoNet [†] +AE	55.89 ± 0.43	73.43 ± 0.35
ProtoNet [†] +AE+PT	55.62 ± 0.45	73.24 ± 0.35
ProtoNet [†] +AE+CL	55.61 ± 0.44	73.52 ± 0.53
CPLAE	56.83 ± 0.44	74.31 ± 0.34

Table 6: Comparison to CPL alternatives on the *miniImageNet* dataset (with Conv4-64 being the backbone).

Method	5-way 1-shot	5-way 5-shot
ProtoNet [†] +AE	55.89 ± 0.43	73.43 ± 0.35
CPLAE (w/o Proto, w/o Proj)	56.18 ± 0.44	73.15 ± 0.35
CPLAE (w/o Proto, w/ Proj)	56.24 ± 0.43	73.35 ± 0.35
CPLAE (w/ Proto, w/o Proj)	56.31 ± 0.44	73.62 ± 0.34
CPLAE (w/ Proto, w/ Proj)	56.83 ± 0.44	74.31 ± 0.34

deformation methods and use their combinations as the alternative data augmentation strategies. Note that the first row of Table 4 involves no data augmentation (i.e., ProtoNet[†]), and the rest results are obtained by ProtoNet[†]+AE. Particularly, the second row means that we only use two deformation methods (i.e., the dimension of augmented embeddings in this case is 3D), and the last two rows use four methods (i.e., the integrated embeddings are of 5D). The comparative results in Table 4 show that the combination of Horizontal Flip, Vertical Flip, and Rotation 270° is the best, and FSL using three deformation methods outperforms FSL using two or four.

Alternative SSL Losses. Our supervised contrastive learning loss L_{cpl} is inspired by the previous self-supervised learning (SSL) and contrastive learning (CL) works. To verify the effectiveness of our CPLAE model with the CPL loss, we compare it to two alternative models with SSL losses. (1) ProtoNet[†]+AE+PT: an SSL loss based on the pretext task (PT) is added into ProtoNet[†]+AE by predicting the augmented embeddings are shuffled or not. (2) ProtoNet[†]+AE+CL: the conventional unsupervised contrastive loss is applied to ProtoNet[†]+AE. Specifically, for each query sample, since we have two augmented embeddings by integration with different orders, we take each one as the anchor in turn. Naturally, the other one is treated as the positive example, while all of the rest are negative examples. The comparative results in Table 5 demonstrate that our proposed CPL is the best choice for FSL.

Alternative Contrastive Learning Losses. The main differences between our proposed contrastive prototype learning (CPL) and the conventional supervised triplet loss [Schroff et al., 2015] (or its improved version N-pair loss [Sohn, 2016]) are in two aspects: (1) Our CPL chooses class pro-

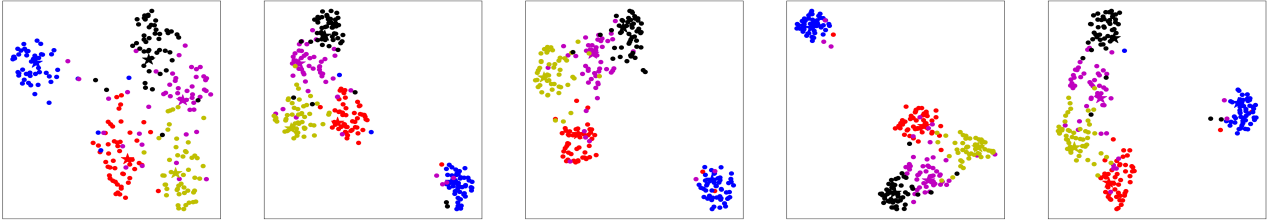


Figure 3: Visualizations of data distributions of the same meta-test episode from the *miniImageNet* dataset using the UMAP algorithm [McInnes et al., 2018] for five FSL models (from left to right): ProtoNet[†], ProtoNet[†]+AE, CPLAE (w/o Proto, w/o Proj, no shuffling), CPLAE (no shuffling), and our CPLAE. Only the 5-way 5-shot setting is considered, and Conv4-64 is used as the feature extractor.

Table 7: Results obtained by varying the number of negative examples for our full CPLAE model on *miniImageNet* (with Conv4-64 as the backbone).

Method	5-way 1-shot	5-way 5-shot
CPLAE ($m = 3$)	56.63 ± 0.44	73.88 ± 0.34
CPLAE ($m = 6$)	56.83 ± 0.44	74.31 ± 0.34
CPLAE ($m = 9$)	56.66 ± 0.43	74.08 ± 0.34
CPLAE ($m = 12$)	56.78 ± 0.44	74.20 ± 0.30
CPLAE ($m = 15$)	56.63 ± 0.44	74.23 ± 0.34

otypes as anchors, but the triplet/N-pair loss takes each sample to be the anchor in turn. (2) Our CPL adopts a projection head which maps the embeddings into a latent space, but such projection is not considered in the triplet/N-pair loss. Therefore, we conduct a group of experiments to find out the contribution of the prototype-based anchors and projection head in Table 6. Methods with ‘Proto’ use prototypes as anchors and those with ‘Proj’ use the projection head when computing L_{cpl} . We can observe that our novel integration of contrastive learning into FSL, i.e., CPLAE (w/ Proto, w/ Proj), achieves the best results.

Results by Varying the Number of Negative Examples. We have mentioned in Section 3.3 of the main paper that for each positive example belonging to class c , we randomly sample m ($m \leq q = 15$) negative examples from each of the other classes. In this section, we thus conduct experiments by selecting m from $\{3, 6, 9, 12, 15\}$ on *miniImageNet* with Conv4-64 as the backbone. As shown in Table 7, we can see that our CPLAE model is insensitive to the number of negative examples used for CPL in general. Specifically, when $m = 6$, the results are slightly better than those of other alternatives. We thus set $m = 6$ in our algorithm for all experiments.

4.4 VISUALIZATION RESULTS

Based on augmented embeddings, Our CPLAE model takes the class prototypes as anchors and the shuffled embeddings of query samples as positive/negative examples. In

Figure 3, we visualize the data distributions of the same meta-test episode obtained by five FSL models: ProtoNet[†], ProtoNet[†]+AE, CPLAE (w/o Proto, w/o Proj, no shuffling), CPLAE (no shuffling), and our CPLAE. We can observe that: (1) CPLAE (no shuffling) leads to better data clustering structure than CPLAE (w/o Proto, w/o Proj, no shuffling), which suggests that the prototype-based anchors and the projection head are effective. (2) CPLAE with shuffling is better than CPLAE (no shuffling), which indicates the necessity of using the shuffling operation for CPL. More visual results can be found in the supplementary material.

5 CONCLUSION

We have proposed a novel contrastive prototype learning with augmented embedding (CPLAE) model to address the lack of training data problem in FSL. Different from existing embedding-based meta-learning methods, we introduce both data augmentation to form an augmented embedding space and a support set prototype centered loss to complement the conventional query centered loss. Extensive experiments on three widely used benchmarks demonstrate that our CPLAE achieves new state-of-the-art. This work shows for the first time that contrastive learning is effective under the supervised and few-shot learning setting.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China (61976220 and 61832017), Beijing Outstanding Young Scientist Program (BJJWZYJH012019100020098), and Open Project Program Foundation of Key Laboratory of Opto-Electronics Information Processing, Chinese Academy of Sciences (OEIP-O-202006).

References

A. Afrasiyabi, J.-F. Lalonde, and C. Gagné. Associative alignment for few-shot image classification. In *ECCV*,

- 2020.
- K. R. Allen, E. Shelhamer, H. Shin, and J. B. Tenenbaum. Infinite mixture prototypes for few-shot learning. In *ICML*, pages 232–241, 2019.
- A. Antoniou and A. Storkey. Assume, augment and learn: Unsupervised few-shot meta-learning via random labels and data augmentation. *preprint arXiv:1902.09884*, 2019.
- L. Bertinetto, J. F. Henriques, P. H. S. Torr, and A. Vedaldi. Meta-learning with differentiable closed-form solvers. In *ICLR*, 2019.
- T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton. A simple framework for contrastive learning of visual representations. *preprint arXiv:2002.05709*, 2020.
- D. L. Davies and D. W. Bouldin. A cluster separation measure. *TPAMI*, 1(2):224–227, 1979.
- C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135, 2017.
- S. Gidaris and N. Komodakis. Generating classification weights with GNN denoising autoencoders for few-shot learning. In *CVPR*, pages 21–30, 2019.
- S. Gidaris, A. Bursuc, N. Komodakis, P. Perez, and M. Cord. Boosting few-shot visual learning with self-supervision. In *ICCV*, pages 8059–8068, 2019.
- Y. Guo and N. Cheung. Attentive weights generation for few shot learning via information maximization. In *CVPR*, pages 13496–13505, 2020.
- B. Hariharan and R. B. Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *ICCV*, pages 3037–3046, 2017.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020.
- K. Hsu, S. Levine, and C. Finn. Unsupervised learning via meta-learning. In *ICLR*, 2019.
- S. Khodadadeh, L. Bölöni, and M. Shah. Unsupervised meta-learning for few-shot image classification. In *NeurIPS*, pages 10132–10142, 2019.
- P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020.
- J. Kim, T. Kim, S. Kim, and C. D. Yoo. Edge-labeling graph neural network for few-shot learning. In *CVPR*, pages 11–20, 2019.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- K. Lee, S. Maji, A. Ravichandran, and S. Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, pages 10657–10665, 2019.
- F. Li, R. Fergus, and P. Perona. A bayesian approach to unsupervised one-shot learning of object categories. In *ICCV*, pages 1134–1141, 2003.
- F. Li, R. Fergus, and P. Perona. One-shot learning of object categories. *TPAMI*, 28(4):594–611, 2006.
- H. Li, D. Eigen, S. Dodge, M. Zeiler, and X. Wang. Finding task-relevant features for few-shot learning by category traversal. In *CVPR*, pages 1–10, 2019a.
- K. Li, Y. Zhang, K. Li, and Y. Fu. Adversarial feature hallucination networks for few-shot learning. In *CVPR*, pages 13467–13476, 2020.
- W. Li, L. Wang, J. Xu, J. Huo, Y. Gao, and J. Luo. Revisiting local descriptor based image-to-class measure for few-shot learning. In *CVPR*, pages 7260–7268, 2019b.
- Z. Lin, M. Feng, C. N. dos Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio. A structured self-attentive sentence embedding. In *ICLR*, 2017.
- B. Liu, Y. Cao, Y. Lin, Q. Li, Z. Zhang, M. Long, and H. Hu. Negative margin matters: Understanding margin in few-shot classification. In *ECCV*, 2020.
- P. Mangla, N. Kumari, A. Sinha, M. Singh, B. Krishnamurthy, and V. N. Balasubramanian. Charting the right manifold: Manifold mixup for few-shot learning. In *WACV*, 2020.
- L. McInnes, J. Healy, and J. Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *preprint arXiv:1802.03426*, 2018.
- C. Medina, A. Devos, and M. Grossglauser. Self-supervised prototypical transfer learning for few-shot classification. In *ICML Workshop on Automated Machine Learning*, 2020.
- T. Munkhdalai and H. Yu. Meta networks. In *ICML*, pages 2554–2563, 2017.
- A. Nichol, J. Achiam, and J. Schulman. On first-order meta-learning algorithms. *preprint arXiv:1803.02999*, 2018.
- B. Oreshkin, P. R. López, and A. Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, pages 721–731, 2018.

- H. Qi, M. Brown, and D. G. Lowe. Low-shot learning with imprinted weights. In *CVPR*, pages 5822–5830, 2018.
- L. Qiao, Y. Shi, J. Li, Y. Wang, T. Huang, and Y. Tian. Transductive episodic-wise adaptive metric for few-shot learning. In *ICCV*, pages 3603–3612, 2019.
- S. Qiao, C. Liu, W. Shen, and A. L. Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, pages 7229–7238, 2018.
- T. Qin, W. Li, Y. Shi, and Y. Gao. Unsupervised few-shot learning via distribution shift-based augmentation. *preprint arXiv:2004.05805*, 2020.
- S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- A. Ravichandran, R. Bhotika, and S. Soatto. Few-shot learning with embedded class models and shot-free meta training. In *ICCV*, pages 331–339, 2019.
- M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018.
- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell. Meta-learning with latent embedding optimization. In *ICLR*, 2019.
- V. G. Satorras and J. B. Estrach. Few-shot learning with graph neural networks. In *ICLR*, 2018.
- F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015.
- C. Simon, P. Koniusz, R. Nock, and M. Harandi. Adaptive subspaces for few-shot learning. In *CVPR*, pages 4136–4145, 2020.
- J. Snell, K. Swersky, and R. S. Zemel. Prototypical networks for few-shot learning. In *NIPS*, pages 4080–4090, 2017.
- K. Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NIPS*, pages 1849–1857, 2016.
- J.-C. Su, S. Maji, and B. Hariharan. When does self-supervision improve few-shot learning? In *ECCV*, 2020.
- Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele. Meta-transfer learning for few-shot learning. In *CVPR*, pages 403–412, 2019.
- F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. S. Torr, and T. M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018.
- Y. Tian, D. Krishnan, and P. Isola. Contrastive multiview coding. *preprint arXiv:1906.05849*, 2019.
- Y. Tian, Y. Wang, D. Krishnan, J. B. Tenenbaum, and P. Isola. Rethinking few-shot image classification: a good embedding is all you need? In *ECCV*, 2020.
- A. van den Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NIPS*, pages 5998–6008, 2017.
- O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NIPS*, pages 3630–3638, 2016.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, CalTech, 2011.
- Z. Wu, Y. Li, L. Guo, and K. Jia. Parn: Position-aware relation networks for few-shot learning. In *ICCV*, pages 6659–6667, 2019.
- C. Xing, N. Rostamzadeh, B. Oreshkin, and P. O. O. Pinheiro. Adaptive cross-modal few-shot learning. In *NeurIPS*, pages 4847–4857, 2019.
- L. Yang, L. Li, Z. Zhang, X. Zhou, E. Zhou, and Y. Liu. DPGN: distribution propagation graph network for few-shot learning. In *CVPR*, pages 13387–13396, 2020.
- H.-J. Ye, H. Hu, D.-C. Zhan, and F. Sha. Few-shot learning via embedding adaptation with set-to-set functions. In *CVPR*, 2020.
- S. W. Yoon, J. Seo, and J. Moon. Tapnet: Neural network augmented with task-adaptive projection for few-shot learning. In *ICML*, pages 7115–7123, 2019.
- C. Zhang, Y. Cai, G. Lin, and C. Shen. Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In *CVPR*, pages 12203–12213, 2020.
- H. Zhang, J. Zhang, and P. Koniusz. Few-shot learning via saliency-guided hallucination of samples. In *CVPR*, pages 2770–2779, 2019.