

# Associative Alignment for Few-shot Image Classification

Arman Afrasiyabi\*, Jean-François Lalonde\*, Christian Gagné\*†

\*Université Laval, †Canada CIFAR AI Chair, Mila  
arman.afraasiyabi.1@ulaval.ca  
{jflalonde,christian.gagne}@gel.ulaval.ca  
<https://lvsn.github.io/associative-alignment/>

**Abstract.** Few-shot image classification aims at training a model from only a few examples for each of the “novel” classes. This paper proposes the idea of associative alignment for leveraging part of the base data by aligning the novel training instances to the closely related ones in the base training set. This expands the size of the effective novel training set by adding extra “related base” instances to the few novel ones, thereby allowing a constructive fine-tuning. We propose two associative alignment strategies: 1) a metric-learning loss for minimizing the distance between related base samples and the centroid of novel instances in the feature space, and 2) a conditional adversarial alignment loss based on the Wasserstein distance. Experiments on four standard datasets and three backbones demonstrate that combining our centroid-based alignment loss results in absolute accuracy improvements of 4.4%, 1.2%, and 6.2% in 5-shot learning over the state of the art for object recognition, fine-grained classification, and cross-domain adaptation, respectively.

**Keywords:** associative alignment, few-shot image classification

## 1 Introduction

Despite recent progress, generalizing on new concepts with little supervision is still a challenge in computer vision. In the context of image classification, few-shot learning aims to obtain a model that can learn to recognize novel image classes when very few training examples are available.

Meta-learning [9, 36, 42, 47] is a possible approach to achieve this, by extracting common knowledge from a large amount of labeled data (the “base” classes) to train a model that can then learn to classify images from “novel” concepts with only a few examples. This is achieved by repeatedly sampling small subsets from the large pool of base images, effectively simulating the few-shot scenario. Standard transfer learning has also been explored as an alternative method [3, 14, 34]. The idea is to pre-train a network on the base samples and then fine-tune the classification layer on the novel examples. Interestingly, Chen *et al.* [3] demonstrated that doing so performs on par with more sophisticated meta-learning strategies. It is, however, necessary to freeze the feature encoder part of the

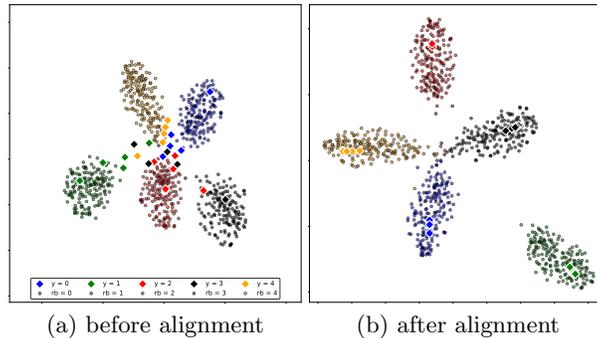


Fig. 1: The use of many related bases (circles) in addition to few novel classes samples (diamonds) allows better discriminative models: (a) using directly related bases may not properly capture the novel classes; while (b) aligning both related base and novel training instances (in the feature space) provides more relevant training data for classification. Plots are generated with t-SNE [30] applied to the ResNet-18 feature embedding before (a) and after (b) the application of the centroid alignment. Points are color-coded by class.

network when fine-tuning on the novel classes since the network otherwise overfits the novel examples. We hypothesize that this hinders performance and that gains could be made if the entire network is adapted to the novel categories.

In this paper, we propose an approach that simultaneously prevents overfitting without restricting the learning capabilities of the network for few-shot image classification. Our approach relies on the standard transfer learning strategy [3] as a starting point, but subsequently exploits base categories that are most similar (in the feature space) to the few novel samples to effectively provide additional training examples. We dub these similar categories the “related base” classes. Of course, the related base classes represent different concepts than the novel classes, so fine-tuning directly on them could confuse the network (see fig. 1-(a)). The key idea of this paper is to *align*, in feature space, the novel examples with the related base samples (fig. 1-(b)).

To this end, we present two possible solutions for associative alignment: by 1) centroid alignment, inspired by ProtoNet [42], benefits from explicitly shrinking the intra-class variations and is more stable to train, but makes the assumption that the class distribution is well-approximated by a single mode. Adversarial alignment, inspired by WGAN [1], does not make that assumption, but its train complexity is greater due to the critic network. We demonstrate, through extensive experiments, that our centroid-based alignment procedure achieves state-of-the-art performance in few-shot classification on several standard benchmarks. Similar results are obtained by our adversarial alignment, which shows the effectiveness of our associative alignment approach.

We present the following contributions. First, we propose two approaches for aligning novel to related base classes in the feature space, allowing for ef-

fective training of entire networks for few-shot image classification. Second, we introduce a strong baseline that combines standard transfer learning [3] with an additive angular margin loss [6], along with early stopping to regularize the network while pre-training on the base categories. We find that this simple baseline actually improves on the state of the art, in the best case by 3% in overall accuracy. Third, we demonstrate through extensive experiments—on four standard datasets and using three well-known backbone feature extractors—that our proposed centroid alignment significantly outperforms the state of the art in three types of scenarios: generic object recognition (gain of 1.7%, 4.4% 2.1% in overall accuracy for 5-shot on *mini*-ImageNet, tieredImageNet and FC100 respectively), fine-grained classification (1.2% on CUB), and cross-domain adaptation (6.2% from *mini*-ImageNet to CUB) using the ResNet-18 backbone.

## 2 Related work

The main few-shot learning approaches can be broadly categorized into meta-learning and standard transfer learning. In addition, data augmentation and regularization techniques (typically in meta-learning) have also been used for few-shot learning. We briefly review relevant works in each category below. Note that several different computer vision problems such as object counting [58], video classification [59], motion prediction [16], and object detection [52] have been framed as few-shot learning. Here, we mainly focus on works from the image classification literature.

*Meta-learning* This family of approaches frames few-shot learning in the form of episodic training [7, 9, 36, 39, 42, 46, 52, 54]. An episode is defined by pretending to be in a few-shot regime while training on the base categories, which are available in large quantities. Initialization- and metric-based approaches are two variations on the episodic training scheme relevant for this work. Initialization-based methods [9, 10, 22] learn an initial model able to adapt to few novel samples with a small number of gradient steps. In contrast, our approach performs a larger number of updates, but requires that the alignment be maintained between the novel samples and their related base examples. Metric-based approaches [2, 12, 21, 25, 27, 33, 42, 44, 45, 47, 53, 57] learn a metric with the intent of reducing the intra-class variations while training on base categories. For example, ProtoNet [42] were proposed to learn a feature space where instances of a given class are located close to the corresponding prototype (centroid), allowing accurate distance-based classification. Our centroid alignment strategy borrows from such distance-based criteria but uses it to match the distributions in the feature space instead of building a classifier.

*Standard transfer learning* The strategy behind this method is to pre-train a network on the base classes and subsequently fine-tune it on the novel examples [3, 14, 34]. Despite its simplicity, Chen *et al.* [3] recently demonstrated that such an approach could result in similar generalization performance compared to meta-learning when deep backbones are employed as feature extractors. However, they

have also shown that the weights of the pre-trained feature extractor must remain frozen while fine-tuning due to the propensity for overfitting. Although the training procedure we are proposing is similar to standard fine-tuning in base categories, our approach allows the training of the entire network, thereby increasing the learned model capacity while improving classification accuracy.

*Regularization trick* Wang *et al.* [51] proposed regression networks for regularization purposes by refining the parameters of the fine-tuning model to be close to the pre-trained model. More recently, Lee *et al.* [24] exploited the implicit differentiation of a linear classifier with hinge loss and  $\mathcal{L}_2$  regularization to the CNN-based feature learner. Dvornik *et al.* [8] uses an ensemble of networks to decrease the classifiers variance.

*Data augmentation* Another family of techniques relies on additional data for training in a few-shot regime, most of the time following a meta-learning training procedure [4, 5, 11, 15, 17, 31, 40, 49, 55, 56]. Several ways of doing so have been proposed, including Feature Hallucination (FH) [17], which learns mappings between examples with an auxiliary generator that then hallucinates extra training examples (in the feature space). Subsequently, Wang *et al.* [49] proposed to use a GAN for the same purpose, and thus address the poor generalization of the FH framework. Unfortunately, it has been shown that this approach suffers from mode collapse [11]. Instead of generating artificial data for augmentation, others have proposed methods to take advantage of additional unlabeled data [13, 37, 26, 50]. Liu *et al.* [29] propose to propagate labels from few labeled data to many unlabeled data, akin to our detection of related bases. We also rely on more data for training, but in contrast to these approaches, our method does not need any new data, nor does it require to generate any. Instead, we exploit the data that is *already available* in the base domain and align the novel domain to the relevant base samples through fine-tuning.

Previous work has also exploited base training data, most related to ours are the works of [4] and [28]. Chen *et al.* [4] propose to use an embedding and deformation sub-networks to leverage additional training samples, whereas we rely on a single feature extractor network which is much simpler to implement and train. Unlike random base example sampling [4] for interpolating novel example deformations in the image space, we propose to borrow the internal distribution structure of the detected related classes in feature space. Besides, our alignment strategies introduce extra criteria to keep the focus of the learner on the novel classes, which prevents the novel classes from becoming outliers. Focused on object detection, Lim *et al.* [28] proposes a model to search similar object categories using a sparse grouped Lasso framework. Unlike [28], we propose and evaluate two associative alignments in the context of few-shot image classification.

From the alignment perspective, our work is related to Jiang *et al.* [20] which stays in the context of zero-shot learning, and proposes a coupled dictionary matching in visual-semantic structures to find matching concepts. In contrast, we propose associative base-novel class alignments along with two strategies for enforcing the unification of the related concepts.

### 3 Preliminaries

Let us assume that we have a large base dataset  $\mathcal{X}^b = \{(\mathbf{x}_i^b, y_i^b)\}_{i=1}^{N^b}$ , where  $\mathbf{x}_i^b \in \mathbb{R}^d$  is the  $i$ -th data instance of the set and  $y_i^b \in \mathcal{Y}^b$  is the corresponding class label. We are also given a small amount of novel class data  $\mathcal{X}^n = \{(\mathbf{x}_i^n, y_i^n)\}_{i=1}^{N^n}$ , with labels  $y_i^n \in \mathcal{Y}^n$  from a set of distinct classes  $\mathcal{Y}^n$ . Few-shot classification aims to train a classifier with only a few examples from each of the novel classes (e.g., 5 or even just 1). In this work, we used the standard transfer learning strategy of Chen *et al.* [3], which is organized into the following two stages.

*Pre-training stage* The learning model is a neural network composed of a feature extractor  $f(\cdot|\theta)$ , parameterized by  $\theta$ , followed by a linear classifier  $c(\mathbf{x}|\mathbf{W}) \equiv \mathbf{W}^\top f(\mathbf{x}|\theta)$ , described by matrix  $\mathbf{W}$ , ending with a scoring function such as softmax to produce the output. The network is trained from scratch on examples from the base categories  $\mathcal{X}^b$ .

*Fine-tuning stage* In order to adapt the network to the novel classes, the network is subsequently fine-tuned on the few examples from  $\mathcal{X}^n$ . Since overfitting is likely to occur if all the network weights are updated, the feature extractor weights  $\theta$  are frozen, with only the classifier weights  $\mathbf{W}$  being updated in this stage.

### 4 Associative alignment

Freezing the feature extractor weights  $\theta$  indeed reduces overfitting, but also limits the learning capacity of the model. In this paper, we strive for the best of both worlds and present an approach which controls overfitting while maintaining the original learning capacity of the model. We borrow the internal distribution structure of a subset of *related* base categories,  $\mathcal{X}^{rb} \subset \mathcal{X}^b$ . To account for the discrepancy between the novel and related base classes, we propose to *align* the novel categories to the related base categories in feature space. Such a mapping allows for a bigger pool of training data while making instances of these two sets more coherent. Note that, as opposed to [4], we do not modify the related base instances in any way: we simply wish to align novel examples to the distributions of their related class instances.

In this section, we first describe how the related base classes are determined. Then, we present our main contribution: the ‘‘centroid associative alignment’’ method, which exploits the related base instances to improve classification performance on novel classes. We conclude by presenting an alternative associative alignment strategy, which relies on an adversarial framework.

#### 4.1 Detecting the related bases

We develop a simple, yet effective procedure to select a set of base categories related to a novel category. Our method associates  $B$  base categories to each novel class. After training  $c(f(\cdot|\theta)|\mathbf{W})$  on  $\mathcal{X}^b$ , we first fine-tune  $c(\cdot|\mathbf{W})$  on  $\mathcal{X}^n$



Fig. 2: Results of related base algorithm in a 5-way 5-shot scenario. Each column represents a different novel class. The top row shows the 5 novel instances, while the bottom row shows 60 randomly selected related base instances with  $B = 10$ .

while keeping  $\theta$  fixed. Then, we define  $\mathbf{M} \in \mathbb{R}^{K^b \times K^n}$  as a base-novel similarity matrix, where  $K^b$  and  $K^n$  are respectively the number of classes in  $\mathcal{X}^b$  and  $\mathcal{X}^n$ . An element  $m_{i,j}$  of the matrix  $\mathbf{M}$  corresponds to the ratio of examples associated to the  $i$ -th base class that are classified as the  $j$ -th novel class:

$$m_{i,j} = \frac{1}{|\mathcal{X}_i^b|} \sum_{(\mathbf{x}_i^b, \cdot) \in \mathcal{X}_i^b} \mathbb{I} \left[ j = \arg \max_{k=1}^{K^n} (c_k(f(\mathbf{x}_i^b | \theta) | \mathbf{W})) \right], \quad (1)$$

where  $c_k(f(\mathbf{x} | \theta) | \mathbf{W})$  is the classifier output  $c(\cdot | \mathbf{W})$  for class  $k$ . Then, the  $B$  base classes with the highest score for a given novel class are kept as the related base for that class. Fig. 2 illustrates example results obtained with this method in a 5-shot, 5-way scenario.

## 4.2 Centroid associative alignment

Let us assume the set of instances  $\mathcal{X}_i^n$  belonging to the  $i$ -th novel class  $i \in \mathcal{Y}^n$ ,  $\mathcal{X}_i^n = \{(\mathbf{x}_j^n, y_j^n) \in \mathcal{X}^n | y_j^n = i\}$ , and the set of related base examples  $\mathcal{X}_i^{rb}$  belonging to the same novel class  $i$  according to the  $g(\cdot | \mathbf{M})$  mapping function,  $\mathcal{X}_i^{rb} = \{(\mathbf{x}_j^b, y_j^b) \in \mathcal{X}^{rb} | g(y_j^b | \mathbf{M}) = i\}$ . The function  $g(y_j^b | \mathbf{M}) : \mathcal{Y}^b \rightarrow \mathcal{Y}^n$  maps base class labels to the novel ones according to the similarity matrix  $\mathbf{M}$ . We wish to find an alignment transformation for matching probability densities  $p(f(\mathbf{x}_{i,k}^n | \theta))$  and  $p(f(\mathbf{x}_{i,l}^{rb} | \theta))$ . Here,  $\mathbf{x}_{i,k}^n$  is the  $k$ -th element from class  $i$  in the novel set, and  $\mathbf{x}_{i,l}^{rb}$  is the  $l$ -th element from class  $i$  in the related base set. This approach has the added benefit of allowing the fine-tuning of all of the model parameters  $\theta$  and  $\mathbf{W}$  with a reduced level of overfitting.

We propose a metric-based centroid distribution alignment strategy. The idea is to enforce intra-class compactness during the alignment process. Specifically, we explicitly push the training examples from the  $i$ -th novel class  $\mathcal{X}_i^n$  towards the centroid of their related examples  $\mathcal{X}_i^{rb}$  in feature space. The centroid  $\boldsymbol{\mu}_i$  of  $\mathcal{X}_i^{rb}$  is computed by

$$\boldsymbol{\mu}_i = \frac{1}{|\mathcal{X}_i^{rb}|} \sum_{(\mathbf{x}_j, \cdot) \in \mathcal{X}_i^{rb}} f(\mathbf{x}_j | \theta), \quad (2)$$

---

**Algorithm 1:**  
 Centroid alignment.
 

---

**Input:** pre-trained model  $c(f(\cdot|\theta)|\mathbf{W})$ ,  
 novel class  $\mathcal{X}^n$ , related base set  $\mathcal{X}^{rb}$ .

**Output:** fine-tuned  $c(f(\cdot|\theta)|\mathbf{W})$ .

**while not done do**

$\tilde{\mathcal{X}}^n \leftarrow$  sample a batch from  $\mathcal{X}^n$   
 $\tilde{\mathcal{X}}^{rb} \leftarrow$  sample a batch from  $\mathcal{X}^{rb}$

evaluate  $\mathcal{L}_{ca}(\tilde{\mathcal{X}}^n, \tilde{\mathcal{X}}^{rb})$ , (eq. 3)  
 $\theta \leftarrow \theta - \eta_{ca} \nabla_{\theta} \mathcal{L}_{ca}(\tilde{\mathcal{X}}^n, \tilde{\mathcal{X}}^{rb})$

evaluate  $\mathcal{L}_{clf}(\tilde{\mathcal{X}}^{rb})$ , (eq. 7)  
 $\mathbf{W} \leftarrow \mathbf{W} - \eta_{clf} \nabla_{\mathbf{W}} \mathcal{L}_{clf}(\tilde{\mathcal{X}}^{rb})$

evaluate  $\mathcal{L}_{clf}(\tilde{\mathcal{X}}^n)$ , (eq. 7)  
 $\mathbf{W} \leftarrow \mathbf{W} - \eta_{clf} \nabla_{\mathbf{W}} \mathcal{L}_{clf}(\tilde{\mathcal{X}}^n)$   
 $\theta \leftarrow \theta - \eta_{clf} \nabla_{\theta} \mathcal{L}_{clf}(\tilde{\mathcal{X}}^n)$

**end**

---

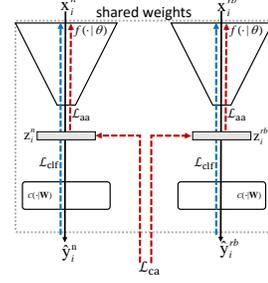


Fig. 3: Schematic overview of our centroid alignment. The feature learner  $f(\cdot|\theta)$  takes an example from novel category  $\mathbf{x}^n$  and an example related base  $\mathbf{x}^{rb}$ . A Euclidean centroid based alignment loss  $\mathcal{L}_{ca}$  (red arrow) aligns the encoded  $\mathbf{x}_i^n$  and  $\mathbf{x}_i^{rb}$ . Blue arrows represent classification loss  $\mathcal{L}_{clf}$ .

where  $N^n$  and  $N^{rb}$  are the number of examples in  $\mathcal{X}^n$  and  $\mathcal{X}^{rb}$ , respectively. This allows the definition of the centroid alignment loss as

$$\mathcal{L}_{ca}(\mathcal{X}^n) = -\frac{1}{N^n N^{rb}} \sum_{i=1}^{K^n} \sum_{(\mathbf{x}_j, \cdot) \in \mathcal{X}^n} \log \frac{\exp[-\|f(\mathbf{x}_j|\theta) - \boldsymbol{\mu}_i\|_2^2]}{\sum_{k=1}^{K^n} \exp[-\|f(\mathbf{x}_j|\theta) - \boldsymbol{\mu}_k\|_2^2]}. \quad (3)$$

Our alignment strategy bears similarities to [42] which also uses eq. 3 in a meta-learning framework. In our case, we use that same equation to match distributions. Fig. 3 illustrates our proposed centroid alignment, and algorithm 1 presents the overall procedure. First, we update the parameters of the feature extraction network  $f(\cdot|\theta)$  using eq. 3. Second, the entire network is updated using a classification loss  $\mathcal{L}_{clf}$  (defined in sec. 5).

### 4.3 Adversarial associative alignment

As an alternative associative alignment strategy, and inspired by WGAN [1], we experiment with training the encoder  $f(\cdot|\theta)$  to perform adversarial alignment using a conditioned critic network  $h(\cdot|\phi)$  based on Wasserstein-1 distance between two probability densities  $p_x$  and  $p_y$ :

$$D(p_x, p_y) = \sup_{\|h\|_L \leq 1} \mathbb{E}_{x \sim p_x}[h(x)] - \mathbb{E}_{x \sim p_y}[h(x)], \quad (4)$$

where sup is the supremum, and  $h$  is a 1-Lipschitz function. Similarly to Arjovsky *et al.* [1], we use a parameterized critic network  $h(\cdot|\phi)$  conditioned by the concatenation of the feature embedding of either  $\mathbf{x}_i^n$  or  $\mathbf{x}_j^{rb}$ , along with the

---

**Algorithm 2:**  
 Adversarial alignment
 

---

**Input:** pre-trained model  $c(f(\cdot|\theta)|\mathbf{W})$ ,  
 novel class  $\mathcal{X}^n$ , related base set  $\mathcal{X}^{rb}$ .

**Output:** fine-tuned  $c(f(\cdot|\theta)|\mathbf{W})$ .

**while not done do**

$\tilde{\mathcal{X}}^n \leftarrow$  sample a batch from  $\mathcal{X}^n$   
      $\tilde{\mathcal{X}}^{rb} \leftarrow$  sample a batch from  $\mathcal{X}^{rb}$

**for**  $i = 0, \dots, n_{\text{critic}}$  **do**

        evaluate  $\mathcal{L}_h(\tilde{\mathcal{X}}^n, \tilde{\mathcal{X}}^{rb})$ , (eq. 5)

        ▷ update critic:

$\phi \leftarrow \phi + \eta_h \nabla_{\phi} \mathcal{L}_h(\tilde{\mathcal{X}}^n, \tilde{\mathcal{X}}^{rb})$

$\phi \leftarrow \text{clip}(\phi, -0.01, 0.01)$

**end**

    evaluate  $\mathcal{L}_{aa}(\tilde{\mathcal{X}}^n)$ , (eq. 6)

$\theta \leftarrow \theta - \eta_{aa} \nabla_{\theta} \mathcal{L}_{aa}(\tilde{\mathcal{X}}^n)$

    evaluate  $\mathcal{L}_{clf}(\tilde{\mathcal{X}}^{rb})$ , (eq. 7)

$\mathbf{W} \leftarrow \mathbf{W} - \eta_{clf} \nabla_{\mathbf{W}} \mathcal{L}_{clf}(\tilde{\mathcal{X}}^{rb})$

    evaluate  $\mathcal{L}_{clf}(\tilde{\mathcal{X}}^n)$ , (eq. 7)

$\mathbf{W} \leftarrow \mathbf{W} - \eta_{clf} \nabla_{\mathbf{W}} \mathcal{L}_{clf}(\tilde{\mathcal{X}}^n)$

$\theta \leftarrow \theta - \eta_{clf} \nabla_{\theta} \mathcal{L}_{clf}(\tilde{\mathcal{X}}^n)$

**end**

---

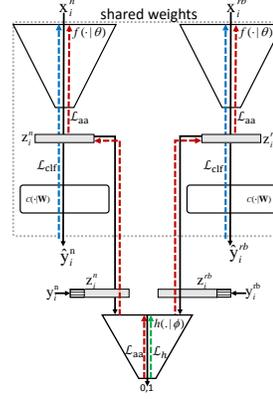


Fig. 4: Overview of our adversarial alignment. The feature learner  $f(\cdot|\theta)$  takes an image  $\mathbf{x}_i^n$  from the  $i$ -th novel class and an example  $\mathbf{x}_i^{rb}$  of the related base. The critic  $h(\cdot|\phi)$  takes the feature vectors and the one-hot class label vector. Green, red and blue arrows present the critic  $\mathcal{L}_h$ , adversarial  $\mathcal{L}_{aa}$  and classification  $\mathcal{L}_{clf}$  losses respectively.

corresponding label  $y_i^n$  encoded as a one-hot vector. Conditioning  $h(\cdot|\phi)$  helps the critic in matching novel categories and their corresponding related base categories. The critic  $h(\cdot|\phi)$  is trained with loss

$$\begin{aligned} \mathcal{L}_h(\mathcal{X}^n, \mathcal{X}^{rb}) &= \frac{1}{N^{rb}} \sum_{(\mathbf{x}_i^{rb}, y_i^{rb}) \in \mathcal{X}^{rb}} h([f(\mathbf{x}_i^{rb}|\theta) y_i^{rb}]|\phi) \\ &\quad - \frac{1}{N^n} \sum_{(\mathbf{x}_i^n, y_i^n) \in \mathcal{X}^n} h([f(\mathbf{x}_i^n|\theta) y_i^n]|\phi), \end{aligned} \quad (5)$$

where,  $[\cdot]$  is the concatenation operator. Then, the encoder parameters  $\theta$  are updated using

$$\mathcal{L}_{aa}(\mathcal{X}^n) = \frac{1}{K^n} \sum_{(\mathbf{x}_i^n, y_i^n) \in \mathcal{X}^n} h([f(\mathbf{x}_i^n|\theta) y_i^n]|\phi). \quad (6)$$

Algorithm 2 summarizes our adversarial alignment method. First, we perform the parameter update of critic  $h(\cdot|\phi)$  using eq. 5. Similar to WGAN [1], we perform  $n_{\text{critic}}$  iterations to optimize  $h$ , before updating  $f(\cdot|\theta)$  using eq. 6. Finally, the entire network is updated by a classification loss  $\mathcal{L}_{clf}$  (defined in sec. 5).

## 5 Establishing a strong baseline

Before evaluating our alignment strategies in sec. 6, we first establish a strong baseline for comparison by following the recent literature. In particular, we build on the work of Chen *et al.* [3] but incorporate a different loss function and episodic early stopping on the pre-training stage.

### 5.1 Classification loss functions

Deng *et al.* [6] have shown that an additive angular margin (“arcmax” hereafter) outperforms other metric learning algorithms for face recognition. The arcmax has a metric learning property since it enforces a geodesic distance margin penalty on the normalized hypersphere, which we think can be beneficial for few-shot classification by helping keep class clusters compact and well-separated.

Let  $\mathbf{z}$  be the representation of  $\mathbf{x}$  in feature space. As per [6], we transform the logit as  $\mathbf{w}_j^\top \mathbf{z} = \|\mathbf{w}_j\| \|\mathbf{z}\| \cos \varphi_j$ , where  $\varphi_j$  is the angle between  $\mathbf{z}$  and  $\mathbf{w}_j$ , the  $j$ -th column in the weight matrix  $\mathbf{W}$ . Each weight  $\|\mathbf{w}_j\| = 1$  by  $l_2$  normalization. Arcmax adds an angular margin  $m$  to the distributed examples on a hypersphere:

$$\mathcal{L}_{\text{clf}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(s \cos(\varphi_{y_i} + m))}{\exp(s \cos(\varphi_{y_i} + m)) + \sum_{\forall j \neq y_i} \exp(s \cos \varphi_j)}, \quad (7)$$

where  $s$  is the radius of the hypersphere on which  $\mathbf{z}$  is distributed,  $N$  the number of examples, and  $m$  and  $s$  are hyperparameters (see sec. 6.1). The overall goal of the margin is to enforce inter-class discrepancy and intra-class compactness.

### 5.2 Episodic early stopping

A fixed number of epochs in the pre-training stage has been commonly used (e.g., [3, 9, 42, 47]), but this might hamper performance in the fine-tuning stage. Using validation error, we observe the necessity of early-stopping in pre-training phase (see supp. mat. for a validation error plot). We thus make the use of episodic early stopping using validation set at pre-training time, specifically by stopping the training when the mean accuracy over a window of recent epochs starts to decrease. The best model in the window is selected as the final result.

## 6 Experimental validation

In the following, we are conducting an experimental evaluation and comparison of the proposed associative alignment strategies for few-shot learning. First, we introduce the datasets used and evaluate the strong baseline from sec. 5.

## 6.1 Datasets and implementation details

*Datasets* We present experiments on four benchmarks: *mini*-ImageNet [47], tieredImageNet [37], and FC100 [33] for generic object recognition; and CUB-200-2011 (CUB) [48] for fine-grained image classification. *mini*-ImageNet is a subset of the ImageNet ILSVRC-12 dataset [38] containing 100 categories and 600 examples per class. We used the same splits as Ravi and Larochelle [36], where 64, 16, and 20 classes are used for the base, validation, and novel classes, respectively. As a larger benchmark, the tieredImageNet [37] is also a subset of ImageNet ILSVRC-12 dataset [38], this time with 351 base, 97 validation, and 160 novel classes respectively. Derived from CIFAR-100 [23], the FC100 dataset [33] contains 100 classes grouped into 20 superclasses to minimize class overlap. Base, validation and novel splits contain 60, 20, 20 classes belonging to 12, 5, and 5 superclasses, respectively. The CUB dataset [48] contains 11,788 images from 200 bird categories. We used the same splits as Hilliard *et al.* [19] using 100, 50, and 50 classes for the base, validation, and novel classes, respectively.

*Network architectures* We experiment with three backbones for the feature learner  $f(\cdot|\theta)$ : 1) a 4-layer convolutional network (“Conv4”) with input image resolution of  $84 \times 84$ , similar to [9, 36, 42]; 2) a ResNet-18 [18] with input size of  $224 \times 224$ ; and 3) a 28-layers Wide Residual Network (“WRN-28-10”) [41] with input size of  $80 \times 80$  in 3 steps of dimension reduction. We use a single hidden layer MLP of 1024 dimensions as the critic network  $h(\cdot|\phi)$  (c.f. sec. 4.3).

*Implementation details* Recall from sec. 3 that training consists of two stages: 1) pre-training using base categories  $\mathcal{X}^b$ ; and 2) fine-tuning on novel categories  $\mathcal{X}^n$ . For pre-training, we use the early stopping algorithm from sec. 5.2 with a window size of 50. Standard data augmentation approaches (i.e., color jitter, random crops, and left-right flips as in [3]) have been employed, and the Adam algorithm with a learning rate of  $10^{-3}$  and batch size of 64 is used for both pre-training and fine-tuning. The arcmax loss (eq. 7) is configured with  $s = 20$  and  $m = 0.1$  which are set by cross validation. In the fine-tuning stage, episodes are defined by randomly selecting  $N = 5$  classes from the novel categories  $\mathcal{X}^n$ .  $k$  examples for each category are subsequently sampled ( $k = 1$  and  $k = 5$  in our experiments). As in Chen *et al.* [3], no standard data augmentation was used in this stage. We used episodic cross-validation to find  $s$  and  $m$  with a fixed encoder. More specifically,  $(s, m)$  were found to be  $(5, 0.1)$  for the Conv4 and  $(5, 0.01)$  for the WRN-28-10 and ResNet-18 backbones. The learning rate for Adam was set to  $10^{-3}$  and  $10^{-5}$  for the centroid and adversarial alignments respectively. Similarly to [1], 5 iterations (inner loop of algorithm 2) were used to train the critic  $h(\cdot|\phi)$ . We fix the number of related base categories as  $B = 10$  (see supp. mat. for an ablation study on  $B$ ). For this reason, we used a relatively large number of categories (50 classes out of the 64 available in *mini*-ImageNet).

Table 1: Preliminary evaluation using *mini*-ImageNet and CUB, presenting 5-way classification accuracy using the Conv4 backbone, with  $\pm$  indicating the 95% confidence intervals over 600 episodes. The best result is boldfaced, while the best result *prior to this work* is highlighted in blue. Throughout this paper, “-” indicates when a paper does not report results in the corresponding scenario.

Method	<i>mini</i> -ImageNet		CUB	
	1-shot	5-shot	1-shot	5-shot
meta learning				
Meta-LSTM [36]	43.44 $\pm$ 0.77	55.31 $\pm$ 0.71	-	-
MatchingNet <sup>‡</sup> [47]	43.56 $\pm$ 0.84	55.31 $\pm$ 0.73	60.52 $\pm$ 0.88	75.29 $\pm$ 0.75
ProtoNet <sup>‡</sup> [42]	49.42 $\pm$ 0.78	68.20 $\pm$ 0.66	51.31 $\pm$ 0.91	70.77 $\pm$ 0.69
MAML <sup>‡</sup> [10]	48.07 $\pm$ 1.75	63.15 $\pm$ 0.91	55.92 $\pm$ 0.95	72.09 $\pm$ 0.76
RelationNet <sup>‡</sup> [44]	50.44 $\pm$ 0.82	65.32 $\pm$ 0.70	<b>62.45 <math>\pm</math> 0.98</b>	76.11 $\pm$ 0.69
tr. learning				
softmax <sup>†</sup>	46.40 $\pm$ 0.72	64.37 $\pm$ 0.59	47.12 $\pm$ 0.74	64.16 $\pm$ 0.71
softmax <sup>†◊</sup>	46.99 $\pm$ 0.73	65.33 $\pm$ 0.60	45.68 $\pm$ 0.86	66.94 $\pm$ 0.84
cosmax <sup>†</sup>	50.92 $\pm$ 0.76	67.29 $\pm$ 0.59	60.53 $\pm$ 0.83	79.34 $\pm$ 0.61
cosmax <sup>†◊</sup>	<b>52.04 <math>\pm</math> 0.82</b>	<b>68.47 <math>\pm</math> 0.60</b>	60.66 $\pm$ 1.04	<b>79.79 <math>\pm</math> 0.75</b>
our baseline (sec. 5)	51.90 $\pm$ 0.79	69.07 $\pm$ 0.62	60.85 $\pm$ 1.07	79.74 $\pm$ 0.64
align.				
adversarial	52.13 $\pm$ 0.99	70.78 $\pm$ 0.60	<b>63.30 <math>\pm</math> 0.94</b>	<b>81.35 <math>\pm</math> 0.67</b>
centroid	<b>53.14 <math>\pm</math> 1.06</b>	<b>71.45 <math>\pm</math> 0.72</b>	62.71 $\pm$ 0.88	80.48 $\pm$ 0.81

<sup>†</sup> our implementation    <sup>◊</sup> with early stopping    <sup>‡</sup> implementation from [3] for CUB

## 6.2 *mini*-ImageNet and CUB with a shallow Conv4 backbone

We first evaluate the new baseline presented in sec. 5 and our associative alignment strategies using the Conv4 backbone on the *mini*-ImageNet (see supp. mat. for evaluations in higher number of ways) and CUB datasets, with corresponding results presented in table 1. We note that arcmax with early stopping improves on using cosmax and softmax with and without early stopping for both the 1- and 5-shot scenarios, on both the *mini*-ImageNet and CUB datasets. We followed the same dataset split configuration, network architecture, and implementation details given in [3] for our testing. Our centroid associative alignment outperforms the state of the art in all the experiments, with gains of 1.24% and 2.38% in 1- and 5-shot over our baseline on *mini*-ImageNet. For CUB, the adversarial alignment provides an additional gain of 0.6% and 0.87% over the centroid one.

## 6.3 *mini*-ImageNet and tieredimageNet with deep backbones

We now evaluate our proposed associative alignment on both the *mini*-ImageNet and tieredimageNet datasets using two deep backbones: ResNet-18 and WRN-28-10. Table 2 compares our proposed alignment methods with several approaches.

*mini-ImageNet* Our centroid associative alignment strategy achieves the best 1- and 5-shot classification tasks on both the ResNet-18 and WRN-28-10 backbones, with notable absolute accuracy improvements of 2.72% and 1.68% over

Table 2: *mini*-ImageNet and tieredImageNet results using ResNet-18 and WRN-28-10 backbones.  $\pm$  denotes the 95% confidence intervals over 600 episodes.

	Method	<i>mini</i> -ImageNet		tieredImageNet	
		1-shot	5-shot	1-shot	5-shot
ResNet-18	TADAM [33]	58.50 $\pm$ 0.30	76.70 $\pm$ 0.30	–	–
	ProtoNet <sup>‡</sup> [42]	54.16 $\pm$ 0.82	73.68 $\pm$ 0.65	61.23 $\pm$ 0.77	80.00 $\pm$ 0.55
	SNAIL [32]	55.71 $\pm$ 0.99	68.88 $\pm$ 0.92	–	–
	IDeMe-Net [4]	59.14 $\pm$ 0.86	74.63 $\pm$ 0.74	–	–
	Activation to Param. [35]	59.60 $\pm$ 0.41	73.74 $\pm$ 0.19	–	–
	MTL [43]	61.20 $\pm$ 1.80	75.50 $\pm$ 0.80	–	–
	TapNet [54]	61.65 $\pm$ 0.15	76.36 $\pm$ 0.10	63.08 $\pm$ 0.15	80.26 $\pm$ 0.12
	VariationalFSL [57]	61.23 $\pm$ 0.26	77.69 $\pm$ 0.17	–	–
	MetaOptNet* [24]	<b>62.64 <math>\pm</math> 0.61</b>	<b>78.63 <math>\pm</math> 0.46</b>	<b>65.99 <math>\pm</math> 0.72</b>	<b>81.56 <math>\pm</math> 0.53</b>
	our baseline (sec. 5)	58.07 $\pm$ 0.82	76.62 $\pm$ 0.58	65.08 $\pm$ 0.19	83.67 $\pm$ 0.51
adversarial alignment	58.84 $\pm$ 0.77	77.92 $\pm$ 0.82	66.44 $\pm$ 0.61	85.12 $\pm$ 0.53	
centroid alignment	59.88 $\pm$ 0.67	<b>80.35 <math>\pm</math> 0.73</b>	<b>69.29 <math>\pm</math> 0.56</b>	<b>85.97 <math>\pm</math> 0.49</b>	
WRN-28-10	LEO [39]	61.76 $\pm$ 0.08	77.59 $\pm$ 0.12	66.33 $\pm$ 0.09	81.44 $\pm$ 0.12
	wDAE [15]	61.07 $\pm$ 0.15	76.75 $\pm$ 0.11	68.18 $\pm$ 0.16	83.09 $\pm$ 0.12
	CC+rot [13]	62.93 $\pm$ 0.45	79.87 $\pm$ 0.33	70.53 $\pm$ 0.51	84.98 $\pm$ 0.36
	Robust-dist++ [39]	63.28 $\pm$ 0.62	<b>81.17 <math>\pm</math> 0.43</b>	–	–
	Transductive-ft [7]	<b>65.73 <math>\pm</math> 0.68</b>	78.40 $\pm$ 0.52	<b>73.34 <math>\pm</math> 0.71</b>	<b>85.50 <math>\pm</math> 0.50</b>
	our baseline (sec. 5)	63.28 $\pm$ 0.71	78.31 $\pm$ 0.57	68.47 $\pm$ 0.86	84.11 $\pm$ 0.65
	adversarial alignment	64.79 $\pm$ 0.93	82.02 $\pm$ 0.88	73.87 $\pm$ 0.76	84.95 $\pm$ 0.59
centroid alignment	<b>65.92 <math>\pm</math> 0.60</b>	<b>82.85 <math>\pm</math> 0.55</b>	<b>74.40 <math>\pm</math> 0.68</b>	<b>86.61 <math>\pm</math> 0.59</b>	

<sup>‡</sup> Results are from [3] for *mini*-ImageNet and from [24] for tieredImageNet, \* ResNet-12

MetaOptNet [24] and Robust-dist++ [8] respectively. The single case where a previous method achieves superior results is that of MetaOptNet, which outperforms our method by 2.76% in 1-shot. For the WRN-28-10 backbone, we achieve similar results to Transductive-ft [7] for 1-shot, but outperform their method by 4.45% in 5-shot. Note that unlike IDeMe-Net [4], SNAIL [32] and TADAM [33], which make use of extra modules, our method achieves significant improvements over these methods without any changes to the backbone.

*tieredImageNet* Table 2 also shows that our centroid associative alignment outperforms the compared methods on tieredImageNet in both 1- and 5-shot scenarios. Notably, our centroid alignment results in a gain of 3.3% and 4.41% over MetaOptNet [24] using the ResNet-18. Likewise, our centroid alignment gains 1.06% and 1.11% over the best of the compared methods using WRN-28-10.

#### 6.4 FC100 and CUB with a ResNet-18 backbone

We present additional results on the FC100 and CUB datasets with a ResNet-18 backbone in table 3. In FC100, our centroid alignment gains 0.73% and 2.14%

Table 3: Results on the FC100 and CUB dataset using ResNet-18 backbones.  $\pm$  denotes the 95% confidence intervals over 600 episodes. The best result is boldfaced, while the best result *prior to this work* is highlighted in blue.

Method	FC100		CUB	
	1-shot	5-shot	1-shot	5-shot
Robust-20 [8]	–	–	58.67 $\pm$ 0.65	75.62 $\pm$ 0.48
GNN-LFT [45]	–	–	51.51 $\pm$ 0.80	73.11 $\pm$ 0.68
RelationNet <sup>‡</sup> [44]	–	–	67.59 $\pm$ 1.02	82.75 $\pm$ 0.58
ProtoNet <sup>‡</sup> [42]	40.5 $\pm$ 0.6	55.3 $\pm$ 0.6	<b>71.88 <math>\pm</math> 0.91</b>	<b>87.42 <math>\pm</math> 0.48</b>
TADAM [33]	40.1 $\pm$ 0.4	56.1 $\pm$ 0.4	–	–
MetaOptNet <sup>‡</sup> [24]	41.1 $\pm$ 0.6	55.5 $\pm$ 0.6	–	–
MTL [43]	<b>45.1 <math>\pm</math> 1.8</b>	<b>57.6 <math>\pm</math> 0.9</b>	–	–
Transductive-ft [7]	43.2 $\pm$ 0.6	<b>57.6 <math>\pm</math> 0.6</b>	–	–
our baseline (sec. 5)	40.84 $\pm$ 0.71	57.02 $\pm$ 0.63	71.71 $\pm$ 0.86	85.74 $\pm$ 0.49
adversarial	43.44 $\pm$ 0.71	58.69 $\pm$ 0.56	70.80 $\pm$ 1.12	88.04 $\pm$ 0.54
centroid	<b>45.83 <math>\pm</math> 0.48</b>	<b>59.74 <math>\pm</math> 0.56</b>	<b>74.22 <math>\pm</math> 1.09</b>	<b>88.65 <math>\pm</math> 0.55</b>

<sup>‡</sup> implementation from [3] for CUB, and from [24] for FC100

over MTL [43] in 1- and 5-shot respectively. We also observe improvements in CUB with our associative alignment approaches, with the centroid alignment outperforming ProtoNet [42] by 2.3% in 1-shot and 1.2% in 5-shot. We outperform Robust-20 [8], an ensemble of 20 networks, by 4.03% and 4.15% on CUB.

## 6.5 Cross-domain evaluation

We also evaluate our alignment strategies in cross-domain image classification. Here, following [3], the base categories are drawn from *mini*-ImageNet, but the novel categories are from CUB. As shown in table 4, we gain 1.3% and 5.4% over the baseline in the 1- and 5-shot, respectively, with our proposed centroid alignment. Adversarial alignment falls below the baseline in 1-shot by -1.2%, but gains 5.9% in 5-shot. Overall, our centroid alignment method shows absolute accuracy improvements over the state of the art (i.e., cosmax [3]) of 3.8% and 6.0% in 1- and 5- shot respectively. We also outperform Robust-20 [8], an ensemble of 20 networks, by 4.65% for 5-shot on *mini*-ImageNet to CUB cross-domain. One could argue that the three bird categories (i.e., house finch, robin, and toucan) in *mini*-ImageNet bias the cross-domain evaluation. Re-training the approach by excluding these classes resulted in a similar performance as shown in table 4.

## 7 Discussion

This paper presents the idea of associative alignment for few-shot image classification, which allows for higher generalization performance by enabling the

Table 4: Cross-domain results from *mini*-ImageNet to CUB in 1-shot, 5-shot, 10-shot scenarios using a ResNet-18 backbone.

Method	1-shot	5-shot	10-shot
ProtoNet <sup>‡</sup> [49]	–	62.02 ± 0.70	–
MAML <sup>‡</sup> [10]	–	51.34 ± 0.72	–
RelationNet <sup>‡</sup> [44]	–	57.71 ± 0.73	–
Diverse 20 [8]	–	<b>66.17 ± 0.73</b>	–
cosmax <sup>†</sup> [3]	<b>43.06 ± 1.01</b>	64.38 ± 0.86	<b>67.56±0.77</b>
our baseline (sec. 5)	45.60 ± 0.94	64.93 ± 0.95	68.95±0.78
adversarial	44.37 ± 0.94	70.80 ± 0.83	79.63 ±0.71
adversarial*	44.65 ± 0.88	71.48 ± 0.96	78.52 ±0.70
centroid	46.85 ± 0.75	70.37 ± 1.02	<b>79.98</b> ±0.80
centroid*	<b>47.25</b> ± 0.76	<b>72.37</b> ± 0.89	79.46 ±0.72

\* without birds (house finch, robin, toucan) in base classes

† our implementation, with early stopping, ‡ implementation from [3]

training of the entire network, still while avoiding overfitting. To do so, we design a procedure to detect related base categories for each novel class. Then, we proposed a centroid-based alignment strategy to keep the intra-class alignment while performing updates for the classification task. We also explored an adversarial alignment strategy as an alternative. Our experiments demonstrate that our approach, specifically the centroid-based alignment, outperforms previous works in almost all scenarios. The current limitations of our work provide interesting future research directions. First, the alignment approach (sec. 4) might include irrelevant examples from the base categories, so using categorical semantic information could help filter out bad samples. An analysis showed that  $\sim 12\%$  of the samples become out-of-distribution (OOD) using a centroid nearest neighbour criteria on *mini*ImageNet in 5-way 1- and 5-shot using ResNet-18. Classification results were not affected significantly by discarding OOD examples at each iteration. Second, the multi-modality of certain base categories look inevitable and might degrade the generalization performance compared to the single-mode case assumed by our centroid alignment strategy. Investigating the use of a mixture family might therefore improve generalization performance. Finally, our algorithms compute the related base once and subsequently keep them fixed during an episode, not taking into account the changes applied to the latent space during the episodic training. Therefore, a more sophisticated dynamic sampling mechanism could be helpful in the finetuning stage.

## Acknowledgement

This project was supported by funding from NSERC-Canada, Mitacs, Prompt-Québec, and E Machine Learning. We thank Ihsen Hedhli, Saed Moradi, Marc-André Gardner, and Annette Schwerdtfeger for proofreading of the manuscript.

## References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein gan. arXiv preprint arXiv:1701.07875 (2017)
2. Bertinetto, L., Henriques, J.F., Torr, P., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. In: The International Conference on Learning Representations (2019)
3. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C.F., Huang, J.B.: A closer look at few-shot classification. arXiv preprint arXiv:1904.04232 (2019)
4. Chen, Z., Fu, Y., Wang, Y.X., Ma, L., Liu, W., Hebert, M.: Image deformation meta-networks for one-shot learning. In: The Conference on Computer Vision and Pattern Recognition (2019)
5. Chu, W.H., Li, Y.J., Chang, J.C., Wang, Y.C.F.: Spot and learn: A maximum-entropy patch sampler for few-shot image classification. In: The Conference on Computer Vision and Pattern Recognition (2019)
6. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: The Conference on Computer Vision and Pattern Recognition (2019)
7. Dhillon, G.S., Chaudhari, P., Ravichandran, A., Soatto, S.: A baseline for few-shot image classification. arXiv preprint arXiv:1909.02729 (2019)
8. Dvornik, N., Schmid, C., Mairal, J.: Diversity with cooperation: Ensemble methods for few-shot classification. In: The International Conference on Computer Vision (2019)
9. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: The International Conference on Machine Learning (2017)
10. Finn, C., Xu, K., Levine, S.: Probabilistic model-agnostic meta-learning. In: Advances in Neural Information Processing Systems (2018)
11. Gao, H., Shou, Z., Zareian, A., Zhang, H., Chang, S.F.: Low-shot learning via covariance-preserving adversarial augmentation networks. In: Advances in Neural Information Processing Systems (2018)
12. Garcia, V., Bruna, J.: Few-shot learning with graph neural networks. arXiv preprint arXiv:1711.04043 (2017)
13. Gidaris, S., Bursuc, A., Komodakis, N., Pérez, P., Cord, M.: Boosting few-shot visual learning with self-supervision. In: The International Conference on Computer Vision (2019)
14. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: The Conference on Computer Vision and Pattern Recognition (2018)
15. Gidaris, S., Komodakis, N.: Generating classification weights with gnn denoising autoencoders for few-shot learning. arXiv preprint arXiv:1905.01102 (2019)
16. Gui, L.Y., Wang, Y.X., Ramanan, D., Moura, J.M.F.: Few-shot human motion prediction via meta-learning. In: The European Conference on Computer Vision (2018)
17. Hariharan, B., Girshick, R.: Low-shot visual recognition by shrinking and hallucinating features. In: The International Conference on Computer Vision (2017)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: The Conference on Computer Vision and Pattern Recognition (2016)
19. Hilliard, N., Phillips, L., Howland, S., Yankov, A., Corley, C.D., Hodas, N.O.: Few-shot learning with metric-agnostic conditional embeddings. arXiv preprint arXiv:1802.04376 (2018)

20. Jiang, H., Wang, R., Shan, S., Chen, X.: Learning class prototypes via structure alignment for zero-shot recognition. In: The European Conference on Computer Vision (2018)
21. Kim, J., Oh, T.H., Lee, S., Pan, F., Kweon, I.S.: Variational prototyping-encoder: One-shot learning with prototypical images. In: The Conference on Computer Vision and Pattern Recognition (2019)
22. Kim, T., Yoon, J., Dia, O., Kim, S., Bengio, Y., Ahn, S.: Bayesian model-agnostic meta-learning. arXiv preprint arXiv:1806.03836 (2018)
23. Krizhevsky, A., Nair, V., Hinton, G.: Cifar-10 and cifar-100 datasets. URL: <https://www.cs.toronto.edu/kriz/cifar.html> (2009)
24. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: The Conference on Computer Vision and Pattern Recognition (2019)
25. 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: The Conference on Computer Vision and Pattern Recognition (2019)
26. Li, X., Sun, Q., Liu, Y., Zhou, Q., Zheng, S., Chua, T.S., Schiele, B.: Learning to self-train for semi-supervised few-shot classification. In: Advances in Neural Information Processing Systems (2019)
27. Lifchitz, Y., Avrithis, Y., Picard, S., Bursuc, A.: Dense classification and implanting for few-shot learning. In: The Conference on Computer Vision and Pattern Recognition (2019)
28. Lim, J.J., Salakhutdinov, R.R., Torralba, A.: Transfer learning by borrowing examples for multiclass object detection. In: Advances in Neural Information Processing Systems (2011)
29. Liu, B., Wu, Z., Hu, H., Lin, S.: Deep metric transfer for label propagation with limited annotated data. In: The IEEE International Conference on Computer Vision (ICCV) Workshops (Oct 2019)
30. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* (2008)
31. Mehrotra, A., Dukkipati, A.: Generative adversarial residual pairwise networks for one shot learning. arXiv preprint arXiv:1703.08033 (2017)
32. Mishra, N., Rohaninejad, M., Chen, X., Abbeel, P.: A simple neural attentive meta-learner. arXiv preprint arXiv:1707.03141 (2017)
33. Oreshkin, B., López, P.R., Lacoste, A.: Tadam: Task dependent adaptive metric for improved few-shot learning. In: Advances in Neural Information Processing Systems (2018)
34. Qi, H., Brown, M., Lowe, D.G.: Low-shot learning with imprinted weights. In: The Conference on Computer Vision and Pattern Recognition (2018)
35. Qiao, S., Liu, C., Shen, W., Yuille, A.L.: Few-shot image recognition by predicting parameters from activations. In: The Conference on Computer Vision and Pattern Recognition (2018)
36. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning (2016)
37. 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. arXiv preprint arXiv:1803.00676 (2018)
38. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *The International Journal of Computer Vision* (2015)

39. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. arXiv preprint arXiv:1807.05960 (2018)
40. Schwartz, E., Karlinsky, L., Shtok, J., Harary, S., Marder, M., Kumar, A., Feris, R., Giryes, R., Bronstein, A.: Delta-encoder: an effective sample synthesis method for few-shot object recognition. In: Advances in Neural Information Processing Systems (2018)
41. Sergey, Z., Nikos, K.: Wide residual networks. In: British Machine Vision Conference (2016)
42. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: Advances in Neural Information Processing Systems (2017)
43. Sun, Q., Liu, Y., Chua, T.S., Schiele, B.: Meta-transfer learning for few-shot learning. In: The Conference on Computer Vision and Pattern Recognition (2019)
44. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: The Conference on Computer Vision and Pattern Recognition (2018)
45. Tseng, H.Y., Lee, H.Y., Huang, J.B., Yang, M.H.: Cross-domain few-shot classification via learned feature-wise transformation. arXiv preprint arXiv:2001.08735 (2020)
46. Vilalta, R., Drissi, Y.: A perspective view and survey of meta-learning. Artificial Intelligence Review (2002)
47. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: Advances in Neural Information Processing Systems (2016)
48. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
49. Wang, Y.X., Girshick, R., Hebert, M., Hariharan, B.: Low-shot learning from imaginary data. In: The Conference on Computer Vision and Pattern Recognition (2018)
50. Wang, Y.X., Hebert, M.: Learning from small sample sets by combining unsupervised meta-training with cnns. In: Advances in Neural Information Processing Systems (2016)
51. Wang, Y.X., Hebert, M.: Learning to learn: Model regression networks for easy small sample learning. In: The European Conference on Computer Vision. Springer (2016)
52. Wang, Y.X., Ramanan, D., Hebert, M.: Meta-learning to detect rare objects. In: The International Conference on Computer Vision (2019)
53. Wertheimer, D., Hariharan, B.: Few-shot learning with localization in realistic settings. In: The Conference on Computer Vision and Pattern Recognition (2019)
54. Yoon, S.W., Seo, J., Moon, J.: Tapnet: Neural network augmented with task-adaptive projection for few-shot learning. arXiv preprint arXiv:1905.06549 (2019)
55. Zhang, H., Zhang, J., Koniusz, P.: Few-shot learning via saliency-guided hallucination of samples. In: The Conference on Computer Vision and Pattern Recognition (2019)
56. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)
57. Zhang, J., Zhao, C., Ni, B., Xu, M., Yang, X.: Variational few-shot learning. In: The International Conference on Computer Vision (2019)
58. Zhao, F., Zhao, J., Yan, S., Feng, J.: Dynamic conditional networks for few-shot learning. In: The European Conference on Computer Vision (2018)
59. Zhu, L., Yang, Y.: Compound memory networks for few-shot video classification. In: The European Conference on Computer Vision (2018)