

---

# Unsupervised Embedding Adaptation via Early-Stage Feature Reconstruction for Few-Shot Classification

---

Dong Hoon Lee<sup>1</sup> Sae-Young Chung<sup>1</sup>

## Abstract

We propose unsupervised embedding adaptation for the downstream few-shot classification task. Based on findings that deep neural networks learn to generalize before memorizing, we develop Early-Stage Feature Reconstruction (ESFR) — a novel adaptation scheme with feature reconstruction and dimensionality-driven early stopping that finds generalizable features. Incorporating ESFR consistently improves the performance of baseline methods on all standard settings, including the recently proposed transductive method. ESFR used in conjunction with the transductive method further achieves state-of-the-art performance on *mini*-ImageNet, *tiered*-ImageNet, and CUB; especially with 1.2%~2.0% improvements in accuracy over the previous best performing method on 1-shot setting.<sup>1</sup>

## 1. Introduction

Deep learning has achieved impressive results on visual recognition tasks. However, it still has difficulty generalizing to novel classes with few examples; while humans can learn to recognize from few experiences. Few-shot classification (Miller et al., 2000; Vinyals et al., 2016; Ravi & Larochelle, 2017) is designed to bridge the gap between the two and has recently attracted substantial attention.

Several works (Liu et al., 2019; Hou et al., 2019; Qiao et al., 2019; Hu et al., 2020; Dhillon et al., 2020; Ziko et al., 2020; Boudiaf et al., 2020) have shown the effectiveness of transductive methods in few-shot classification, showing a significant improvement over inductive methods. While test samples are inaccessible in an inductive few-shot classification setting, one can utilize all the unlabeled test samples together to make an inference in a transductive

setting. The co-existence of labeled- and unlabeled-data in this setting motivates the use of transductive inference or semi-supervised learning. A popular transductive approach is pseudo-label-based methods that progressively update the labels or inference models by predicted test samples. For instance, Liu et al. (2019); Kim et al. (2019) use a neighbor graph for label propagation, Hou et al. (2019); Liu et al. (2020) use predicted labels on test samples to update the class prototypes, and Antoniou & Storkey (2019); Hu et al. (2020) use prediction to produce an intrinsic loss or synthetic gradient. Another line of works utilizes regularization terms on unlabeled test samples. To list a few, Dhillon et al. (2020); Boudiaf et al. (2020) use entropy minimization of the prediction on unlabeled data, Ziko et al. (2020) uses the Laplacian regularization term for graph clustering. These methods are mostly originated from semi-supervised learning studies; approaches in semi-supervised learning are still motivating few-shot classification research.

On the other hand, semi-supervised learning research has recently benefited from unsupervised learning. Rapidly advancing self-supervised learning methods (Caron et al., 2020; Grill et al., 2020; Chen et al., 2020) have shown strong performance on semi-supervised image classification tasks. A popular approach is to use a self-supervision loss for representation learning to acquire more general features. Learned representations are then used with fine-tuning (Caron et al., 2020; Grill et al., 2020; Chen et al., 2020) or other semi-supervised learning methods (Zhai et al., 2019; Kim et al., 2021) for downstream tasks. These studies achieved state-of-the-art performance on semi-supervised learning tasks, especially in settings with extremely few labels.

The success of unsupervised learning on semi-supervised tasks suggests the potential benefit of finding shared features or patterns without labels in relevant research areas. In few-shot classification, several works (Gidaris et al., 2019; Su et al., 2020) use additional self-supervision loss during the training of base datasets to learn more general features. However, the use of unsupervised learning on unlabeled data that appears in test-time is less studied. In this work, we study unsupervised learning for adaptation to satisfy the thirst.

---

<sup>1</sup>School of Electrical Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea. Correspondence to: Dong Hoon Lee <donghoonlee@kaist.ac.kr>.

Our contributions are summarized as follows:

- We find that early generalized features during unsupervised training are valuable for recognizing novel classes of few-shot classification. Based on recent studies of deep neural network’s training dynamics, we explain the finding with experiments. (Section 3.1)
- Based on the finding, we construct a novel embedding adaptation scheme with (1) feature reconstruction training and (2) dimensionality-driven early stopping. Our method provides task-adapted embeddings composed of desirable-shared features, which are more likely to be task-relevant and valuable for the few-shot classification. Our method is used as a plug-and-play module for few-shot methods. (Section 3.3)
- We test our method, ESFR, used in conjunction with baseline methods in the standard few-shot classification benchmarks. ESFR consistently improves the performance of baselines; adding ESFR to the transductive method achieves the state-of-the-art performance on *mini-ImageNet*, *tiered-ImageNet*, and *CUB*. Particularly in the scarce-label setting of 1-shot, our method outperforms the previous state-of-the-art with accuracies of 1.2%~2.0%. (Section 5.3)

## 2. Preliminaries

### 2.1. Problem Setting

In a few-shot classification task, a small labeled support set  $S = \{(x_s^i, y_s^i)\}_{i=1}^{|S|}$  and unlabeled query set  $Q = \{(x_q^i, \cdot)\}_{i=1}^{|Q|}$  are given.<sup>2</sup> Both support samples and query samples are from the same novel classes that are never seen during training. The goal of few-shot classification is to classify query (test) samples by few examples in the support set. In the usual setting, the support set has  $K = 1$  or 5 examples per  $N = 5$  novel classes, and we call this an  $N$ -way  $K$ -shot problem.

We address a transductive few-shot classification task where all query samples are accessible. Since learning from few samples without prior is extremely hard, we use a pre-trained embedding network  $f$  that is trained on the base dataset as in [Rusu et al. \(2019\)](#); [Wang et al. \(2019\)](#); [Hu et al. \(2020\)](#); [Ziko et al. \(2020\)](#). We denote  $S_f = \{(f(x_s^i), y_s^i)\}_{i=1}^{|S|}$  and  $Q_f = \{(f(x_q^i), \cdot)\}_{i=1}^{|Q|}$  as the support set and the query set in the embedding domain, respectively.

Our interest is in task-adapted embeddings (representations) that are useful in the given few-shot task. We construct the embeddings with a module  $g_\phi$  on top of  $f$  by training on the union of the support set and query set.

<sup>2</sup> $(x, y)$  denotes an image and its label.

### 2.2. Local Intrinsic Dimensionality (LID)

We briefly explain LID that our method uses as early stopping criterion. LID is a statistical version of an expansion-based Intrinsic Dimension (ID) that provides an estimated subspace dimension of local regions. Recently, substantial redundant dimensions of modern deep neural networks lead to the wide use of ID and LID to track and analyze training ([Amsaleg et al., 2017](#); [Ma et al., 2018a;b](#); [Ansuini et al., 2019](#); [Gong et al., 2019](#)). The formal definition of LID is given as ([Amsaleg et al., 2015](#); [Houle, 2017a;b](#)):

**Definition 1** (Local Intrinsic Dimensionality). *Given a data point  $x$ , let  $r > 0$  be a continuous random distance variable from  $x$ . For the cumulative density function  $F_x(r)$ , the LID of  $x$  at distance  $r$  is defined as:*

$$\text{LID}(r; F_x) \stackrel{\text{def}}{=} \lim_{\epsilon \rightarrow 0^+} \frac{\ln F_x((1 + \epsilon)r) - \ln F_x(r)}{\ln(1 + \epsilon)}, \quad (1)$$

whenever the limit exists. The LID at  $x$  is then defined as the limit of distance  $r \rightarrow 0^+$ :

$$\text{LID}(x) \stackrel{\text{def}}{=} \lim_{r \rightarrow 0^+} \text{LID}(r; F_x). \quad (2)$$

Since the density function of the distance variable is usually unknown, the exact value of LID is hard to acquire. We use the maximum likelihood estimation by [Amsaleg et al. \(2015\)](#) to calculate the LID estimates as follows:

$$\widehat{\text{LID}}(x) = - \left[ \frac{1}{m} \sum_{i=1}^m \ln \frac{r_i(x)}{r_m(x)} \right]^{-1}, \quad (3)$$

where  $r_i(x)$  indicates the distance<sup>3</sup> between  $x$  and its  $i$ -th nearest neighbor. The number of the nearest neighbor  $m$  should be chosen appropriately to make estimation local but stabilized.<sup>4</sup>

We refer to [Amsaleg et al. \(2015\)](#); [Houle \(2017a;b\)](#) for more details about LID and its estimation methods.

## 3. Methodology

Figure 1 illustrates the usage and overview of our method. As a plug-and-play module, our method provides task-adapted embeddings to other few-shot classification methods. Our method is mainly composed of feature reconstruction training and LID-based early stopping. We will explain feature reconstruction training (Section 3.1); LID-based early stopping (Section 3.2); the overall method (Section 3.3), in the following.

<sup>3</sup>We use Euclidean distance.

<sup>4</sup>We set  $m = 20$  throughout experiments as in [Ma et al. \(2018a\)](#).

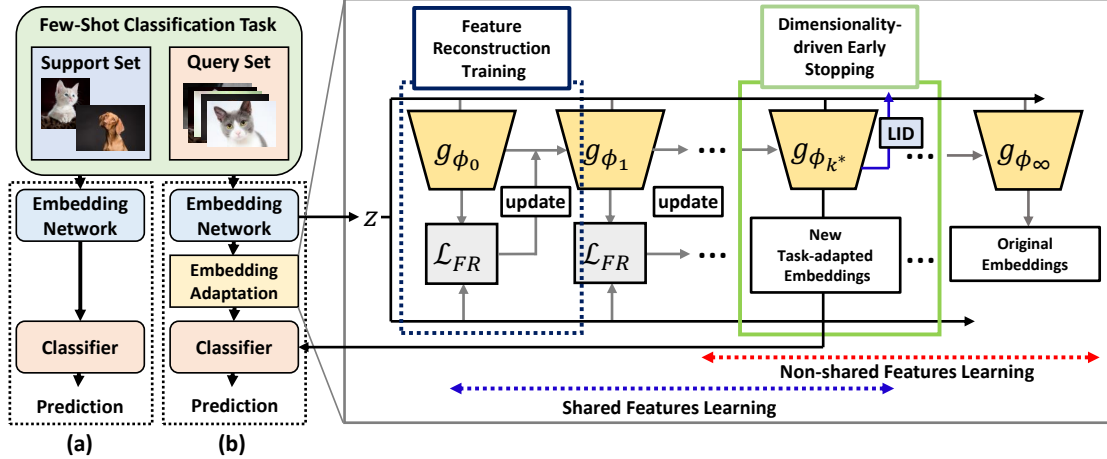


Figure 1. Overview of our method. 1-(a) shows the case without embedding adaptation, and 1-(b) shows the case with embedding adaptation. Our scheme mainly consists of *feature reconstruction training* and *dimensionality-driven early stopping*, and provides new embeddings of generalizable features for the downstream few-shot task.

### 3.1. Feature Reconstruction

Our main idea is based on the findings that *deep neural networks learn to generalize before memorizing* to abstract task-useful features. We design feature-level reconstruction training, which is unsupervised learning and builds on prior knowledge given by embeddings. As mentioned before, unsupervised learning for few-shot adaptation tends to have less attention; and we find that naively applied unsupervised learning for few-shot adaptation often fails.<sup>5</sup> Moreover, the behavior of unsupervised learning with a pre-trained embedding network is often ambiguous. For instance, contrastive learning heavily relies on augmentations, while augmentations affect embeddings differently depending on how the embedding network is pre-trained. Instead, we find that our feature reconstruction training can be used to adapt embeddings for few-shot classification.

We explain our feature reconstruction training. For a few-shot classification task with embedding support set  $S_f$  and query set  $Q_f$ ; we train a reconstruction module  $g_\phi$  using the following feature level reconstruction loss:

$$\mathcal{L}(\phi) = \frac{1}{|S_f \cup Q_f|} \sum_{z \in S_f \cup Q_f} d_{\cos}(z, g_\phi(z)), \quad (4)$$

where  $d_{\cos}$  denotes the cosine distance. Both  $z$  and  $g_\phi(z)$  are preprocessed<sup>6</sup> embeddings, but their expressions are omitted for notational simplicity. We note that for a newly given few-shot classification task, the weight  $\phi$  of the reconstruction module is randomly re-initialized.

We investigate the behavior of reconstruction modules during feature reconstruction training *w.r.t.* the downstream

<sup>5</sup>We test self-supervised learning models of rotation (Gidaris et al., 2018) and jigsaw (Noroozi & Favaro, 2016) in Appendix E.

<sup>6</sup>Details are described in the experiment section.

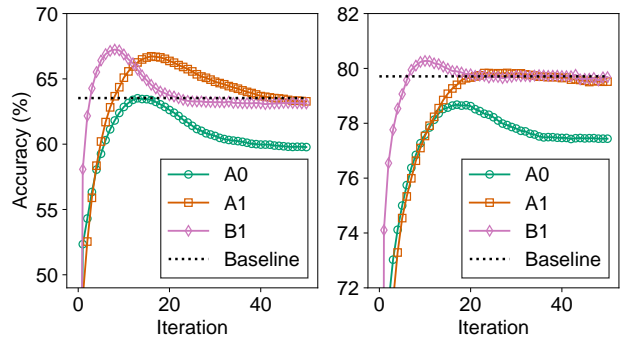


Figure 2. Few-shot classification accuracy with **A0**, **A1**, and **B1** during feature reconstruction training. We use ResNet-18 backbone, *mini-ImageNet* dataset, and nearest neighbor classifier. For  $g_{\phi_1}$ , we use a 4-layer neural network with 256-128-256-512 units. Other settings are described in the experiment section. (left) shows the accuracy on 1-shot and (right) shows the accuracy on 5-shot setting.

few-shot task. We train two types of reconstruction modules,  $g_{\phi_1}$  with compressed (bottleneck) hidden layer as in conventional auto-encoders, and  $g_{\phi_2}$ ; without compression. We evaluate few-shot classification accuracy with three embeddings **A0**, **A1**, **B1**; where **A0** is the middle compressed hidden layer output of  $g_{\phi_1}$ ; **A1** and **B1** are the reconstructed output of  $g_{\phi_1}$  and  $g_{\phi_2}$ , respectively.

Figure 2 shows an interesting behavior that few-shot classification accuracies, with embeddings of reconstruction modules, initially increase then decrease. Moreover, the peak accuracy of **B1** exceeds the baseline accuracy of the original embedding, on both 1- and 5-shot settings. We believe that recent studies of Deep Neural Networks (DNNs) explain such behavior. Several works (Arpit et al., 2017; Lampinen & Ganguli, 2019; Stephenson et al., 2021) observe a property that DNNs learn to generalize before memorizing. To be more specific, Arpit et al. (2017) reports that DNNs learn

patterns first before memorization, with experiments on the mixture of well-structured real data and noisy data. Further research Lampinen & Ganguli (2019); Stephenson et al. (2021) provide analytical explanations on the property that learning speeds between generalization of patterns and memorization of noise are different.<sup>7</sup> We argue that the behavior, shown in Figure 2, is a result of the DNNs’ property. By the property, reconstruction modules learn shared features faster since they form certain patterns or correlations among data, while non-shared features are learned later since they are less generalizable. Generalizable shared features are more likely to be task-relevant in classification; the difference in learning speeds between shared features and non-shared features explains the initial increase of accuracies.

Our main idea is to use the behavior shown in Figure 2 for embedding adaptation. The behavior provides an opportunity to achieve improved few-shot classification performance by acquiring new embeddings of generalized features. To do this, we find that non-compressed and non-encoded embedding (**B1**) performs the best. Throughout the rest of our work, we use the reconstructed output  $g_\phi$  that is non-compressed as task-adapted embeddings.




Figure 3. The figure shows the accuracy and the LID curve during feature reconstruction training, with (Acc:dropout, LID:dropout) and without (Acc, LID) dropout perturbation. We use the same setting as in Figure 2. The vertical lines indicate the maximum accuracy (blue) and the minimum LID (red) points; for 1-shot, they are overlapped. (left) shows the results in the 1-shot and (right) shows the results in the 5-shot setting.

We propose to perturb the training to further discard the non-shared features. Perturbation with noise does not form generalizable patterns or correlations; thus, perturbation tends to make training harder for less generalizable non-shared features while shared features are still learnable by pattern learning. While additive perturbation is hard to design, due to the unknown distribution of  $z$ , we find that multiplicative perturbation by dropout (Srivastava et al., 2014) is well suited for our purpose. Our feature reconstruction training

with input dropout uses the following loss function<sup>8</sup>:

$$\mathcal{L}_{\text{FR}}(\phi) = \frac{1}{|S_f \cup Q_f|} \sum_{z \in S_f \cup Q_f} \mathbb{E}_\mu [d_{\cos}(z, g_\phi(z \circ \mu))], \quad (5)$$

where  $\mu$  is a multiplicative noise implemented with dropout. Figure 3 shows the effect of dropout on the training curve; we can observe that the dropout perturbation results in higher peak accuracies.

### 3.2. Dimensionality Driven Early Stopping

Utilization of early retained generalizable features seems a promising idea for the downstream few-shot classification. However, determining the optimal early stopping time is not straightforward. Here, we suggest using Local Intrinsic Dimensionality (LID) as an early stopping criterion. Recent work by Ma et al. (2018b) used LID to monitor the internal generalization and memorization during training. They argued that LID tends to decrease while DNNs learn to generalize, due to summarizing effect; LID tends to increase while memorizing, as DNNs try to encode detailed information individually. We observe the tendency in our reconstruction training, and we find that LID can be used as an early stopping criterion.

For a given few-shot classification task with embedding support set  $S_f$  and query set  $Q_f$ , we use estimated LID, with reconstruction module  $g_\phi$ , given as:

$$\begin{aligned} \widehat{\text{LID}}(\phi) &= \sum_{z \in S_f \cup Q_f} \widehat{\text{LID}}(g_\phi^{L-2}(z)) \\ &= - \sum_{z \in S_f \cup Q_f} \left[ \frac{1}{m} \sum_{i=1}^m \ln \frac{r_i(g_\phi^{L-2}(z))}{r_m(g_\phi^{L-2}(z))} \right]^{-1}, \quad (6) \end{aligned}$$

where  $g_\phi^{L-2}$  is the hidden representation of the second-to-last layer of  $g_\phi$ ;  $r_i(g_\phi^{L-2}(z))$  denotes the Euclidean distance between  $g_\phi^{L-2}(z)$  and its  $i$ -th nearest neighbor in  $\{g_\phi^{L-2}(z') | z' \in S_f \cup Q_f\}$ . Our  $\widehat{\text{LID}}(\phi)$  is a proxy for the hidden layer subspace dimensionality of the reconstruction module.

In Figure 3, we empirically investigate the relationship between the LID and accuracy during reconstruction training. The result shows that the change of LID can be used to find the early stopping time of the best possible new embeddings; when the LID becomes the lowest or starts to increase. To be more specific, in 1-shot settings, we observe that the LID behaves exactly the opposite of the accuracy curve, for both with and without dropout. For 5-shot settings, the LID behaves almost the opposite of the accuracy curve; however, it has a small misalignment, especially for the case

<sup>7</sup>Generalization is faster than memorization.

<sup>8</sup> $\circ$  indicates the element-wise product.

**Algorithm 1** ESFR

---

**Input:** embedding support set  $S_f$ , embedding query set  $Q_f$ , and few-shot classifier Alg :  $S_f, Q_f \rightarrow \hat{Y}_Q$

**Initialize:**  $\phi^{i=1:N_e}$

**for**  $i = 1$  **to**  $N_e$  **do**

prev\_lid =  $\widehat{\text{LID}}(\phi_0^i)$

**Initialize:** optimizer

**for**  $j = 0$  **to** MAX.ITERATION **do**

$\phi_{j+1}^i \leftarrow \phi_j^i - \nabla_{\phi_j^i} \mathcal{L}(\phi_j^i)$  from equation 5 or 7

lid =  $\widehat{\text{LID}}(\phi_{j+1}^i)$

**if** lid > prev\_lid **then**

$\phi_*^i = \phi_{j+1}^i$

break

**end if**

prev\_lid = lid

**end for**

**end for**

$S^{\text{ESFR}} = \{(z', y) | z' = \frac{1}{N_e} \sum_{i=1}^{N_e} g_{\phi_*^i}(z), (z, y) \in S_f\}$

$Q^{\text{ESFR}} = \{z' | z' = \frac{1}{N_e} \sum_{i=1}^{N_e} g_{\phi_*^i}(z), z \in Q_f\}$

**Output:**  $\hat{Y}_Q = \text{Alg}(S^{\text{ESFR}}, Q^{\text{ESFR}})$

---

with dropout. This seems reasonable since the classifier can further filter out non-shared features by multiple examples of novel classes in 5-shot settings; hence acquiring more shared features at the cost of learning non-shared ones can be advantageous.

### 3.3. Proposed Methods

**ESFR:** We describe our unsupervised adaptation scheme for few-shot classification: Early-Stage Feature Reconstruction (ESFR). Given a few-shot classification task and embedding network, we run feature reconstruction training by  $\mathcal{L}_{\text{FR}}$  (5) with initialized  $\phi$ . At every training iteration, we measure the LID of  $g_\phi$  with task samples by (6) and we early stop the training when the LID starts to increase. Task-adapted embeddings of  $g_\phi(z)$  are used for the classification of the given few-shot task. A wide range of metric-based and fine-tuning few-shot methods can be used with our embeddings for such classification. To reduce the variance by random initial weights of  $\phi$ , we train  $N_e$  reconstruction modules, separately, with different initial weights, and take the center of each sample’s reconstructed embeddings.

**ESFR-Semi:** We further investigate the semi-supervised version of our scheme, ESFR-Semi, by adding the support classification loss to the reconstruction loss (5):

$$\begin{aligned}
 C^j(z, \phi, W, b) &= \text{softmax}_j[Wg_\phi(z) + b] \\
 \mathcal{L}_{\text{CE}}(\phi, W, b) &= -\frac{1}{|S_f|} \sum_{(z_i, y_i) \in S_f} \log C^{y_i}(z_i, \phi, W, b) \\
 \mathcal{L}_{\text{Semi}}(\phi, W, b) &= \mathcal{L}_{\text{FR}}(\phi) + \lambda \mathcal{L}_{\text{CE}}(\phi, W, b) \quad (7)
 \end{aligned}$$

where  $\mathcal{L}_{\text{CE}}$  is the cross-entropy loss given by an affine classifier on new embeddings, and  $\lambda$  is a trade-off parameter. Additional weights  $W$  and  $b$  are jointly trained with  $\phi$ . The trade-off parameter  $\lambda$  is tuned using few-shot classification tasks from validation datasets as in Ziko et al. (2020). Note that the only ESFR-Semi requires (while ESFR does not) few-shot classification task experiences to determine a certain parameter.

The overall methods are described in Algorithm 1.

## 4. Related Work

**Few-Shot Classification (FSC):** Few-shot learning or few-shot classification methods have broad categories: optimization-based methods (Ravi & Larochelle, 2017; Finn et al., 2017; Rusu et al., 2019), distance-based approaches (Vinyals et al., 2016; Snell et al., 2017; Oreshkin et al., 2018), fine-tunings (Wang et al., 2019; Tian et al., 2020; Dhillon et al., 2020), etc. A large portion of these methods is based on meta-learning (Thrun & Pratt, 2012). In meta-learning, training is done in a series of few-shot classification tasks (a.k.a. episodic training) to train the model in a way that reflects test-time scenarios. Several recent studies (Wang et al., 2019; Tian et al., 2020; Dhillon et al., 2020; Ziko et al., 2020; Boudiaf et al., 2020) have questioned the necessity of meta-learning on few-shot classification, reporting competitive performance on few-shot benchmarks without neither episodic training nor few-shot task experiences. These methods solve the few-shot task by fine-tuning<sup>9</sup> a pre-trained embedding network trained on the base dataset with standard cross-entropy loss. Our method lies in this line of research that doesn’t require meta-learning. It seems possible to merge our method with meta-learning; we leave it as future work.

**Unsupervised adaptation for FSC:** We mentioned in the introduction that transductive methods, based on semi-supervised learning techniques, have been widely studied in few-shot classification. In contrast, unsupervised adaptation for few-shot classification is an open field. To the best of our knowledge, we are the first to propose deep unsupervised adaptation for few-shot classification. The closest related works are Rodríguez et al. (2020) and Lightenstein et al. (2020). Embedding Propagation proposed by Rodríguez et al. (2020) iteratively updates embeddings by a linear combination with the nearest neighbors’ embeddings and improves few-shot classification performance. Lightenstein et al. (2020) uses the principal component analysis to acquire major components from given task embeddings; this classical approach shows impressive performance gain. However, these works are limited to an affine transformation

<sup>9</sup>Including the only change of class prototypes or last layer parameters.

of embeddings and do not benefit from deep learning; while our method finds non-linear patterns and correlations by the benefit of deep unsupervised learning. Though unsupervised adaptation in few-shot classification had less attention, our experimental results show that it can outperform conventional transductive baselines, especially when extremely few labeled samples are available in the 1-shot setting.

**Training behavior of DNNs:** Our method is based on the property of Deep Neural Networks’ (DNNs) training behavior that *DNNs learn to generalize before memorizing*. The property is reported and studied in recent works (Arpit et al., 2017; Lampinen & Ganguli, 2019; Stephenson et al., 2021), as mentioned in Section 3.1. Similar to our work, several methods of training DNNs with noisy labels are based on different training behaviors between generalization and memorization. After the report of the property by Arpit et al. (2017), several works (Hendrycks et al., 2019; Oymak et al., 2019; Song et al., 2020) suggest using early stopping to discard noisy information before memorization, Jiang et al. (2018); Yu et al. (2019); Sugiyama (2018) propose to gather the clean samples that exhibit high confident prediction in the early pattern learning for further training. Ma et al. (2018b) offer to use LID to detect and correct noisy label memorization, which strongly motivates our methodology. However, these works are limited to the domain of learning from noisy labels and supervised learning. In our work, we find the behavior also appears in unsupervised adaptation for few-shot classification.

## 5. Experiments

### 5.1. Experimental Settings

**Datasets:** We evaluate our method on three standard datasets of few-shot classification: (1) *mini-ImageNet* (Vinyals et al., 2016) dataset as in Ravi & Larochelle (2017), (2) *tiered-ImageNet* dataset as in Ren et al. (2018), and (3) Caltech-UCSD Birds 200 (CUB) (Welinder et al., 2010) as in Chen et al. (2019). Each dataset is divided into train/val/test splits according to references. All the images are resized to  $84 \times 84$ .

**Baseline Methods:** We investigate three baseline few-shot classification methods (classifiers) used in conjunction with our method.

- *Linear*: Train an affine classifier on labeled support set with embeddings and use the trained classifier to classify the query samples.
- *Nearest Neighbor* (NN): Compute the class prototype by the centroid of support sample embeddings in each class and classify the query sample to the class of the nearest prototype.<sup>10</sup>

<sup>10</sup>We use Euclidean as a distance metric.

- *BD-CSPN*: BD-CSPN (Liu et al., 2020) is chosen as a baseline transductive method. It can be used with pre-trained embeddings and achieves state-of-the-art performance on 5-shot (Table 2). To briefly summarize, BD-CSPN consists of two components: (1) *shifting-term* for removing the cross-class bias between the support set and query set, (2) *Prototype Rectification* (PR) that updates class prototypes by accounting pseudo-labeled query samples. Our particular interest is in PR since it is a baseline approach of transductive methods; a similar scheme was suggested in Ren et al. (2018); Hou et al. (2019).

**Evaluation Protocol:** We evaluate our method on standard 5-way 1- and 5-shot settings with 15 query samples per class. We sampled 2,000 tasks for each experimental result. Our method uses the same fixed hyper-parameters for all experiments and settings. For the semi-supervised version of our method, with trade-off parameter  $\lambda$ , we tune  $\lambda$  by selecting the best among  $\lambda \in [0, 0.1, 0.2, 0.4, 0.8, 1.6]$  from 600 sampled few-shot tasks of validation dataset.

### 5.2. Implementation Details

**Embedding Network:** We follow the embedding network training procedure from Ziko et al. (2020), using backbone architectures: ResNet-18, WRN-28-10. We use *from input to the average-pooled last residual block output* as an embedding network. Embedding networks are trained for 90 epochs, using stochastic gradient descent to minimize the standard cross-entropy on labeled base datasets. Label smoothing with parameter 0.1 is used for more general features. Random cropping, color jittering, and random horizontal flipping are applied for data augmentation.<sup>11</sup> The initial learning rate is 0.1 and shrank by  $\frac{1}{10}$  at 45 and 66 epochs. Mini-batch sizes of 256 and 128 are used for ResNet-18 and WRN-28-10 training, respectively. The best performing embedding network on the 5-way 1-shot task of validation split dataset is chosen; while using the nearest neighbor classifier and l2-normalized embeddings for classification.

**Reconstruction Training:** For the reconstruction module, we use 4-layer fully connected deep neural network with ReLU activation. Each layer has the size of units equal to the embedding dimension, and weight initialization follows the TensorFlow (Abadi et al., 2016) default setting (GlorotUniform). As an optimizer, we use Adam (Kingma & Ba, 2015) with a default learning rate of 1e-3. For dropout (Srivastava et al., 2014), we use the maximum possible rate of 0.5. Finally, we take the ensemble of  $N_e = 5$  reconstruction modules.

**Preprocessing:** Wang et al. (2019) reports the importance of preprocessing in few-shot classification. We apply center-

<sup>11</sup>Data augmentations are used only for embedding network pre-training.

Table 1. Improvement by incorporating our method into baseline methods with ResNet-18/WRN-28-10 backbone on *mini-ImageNet* and *tiered-ImageNet*. † indicates the use of shifting-term (8) during preprocessing.

		ResNet-18				WRN-28-10			
Method		<i>mini-ImageNet</i>		<i>tiered-ImageNet</i>		<i>mini-ImageNet</i>		<i>tiered-ImageNet</i>	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
(i)	Linear	62.45	79.32	68.49	83.77	64.53	80.81	69.78	84.91
	+ ESFR	<b>70.38+7.93</b>	<b>81.6+2.28</b>	<b>76.98+8.49</b>	<b>86.09+2.32</b>	<b>73.33+8.8</b>	<b>83.65+2.84</b>	<b>78.57+8.79</b>	<b>87.37+2.46</b>
(ii)	NN	64.04	79.71	71.60	84.62	66.73	81.85	72.97	85.74
	+ ESFR	<b>70.94+6.9</b>	<b>81.61+1.9</b>	<b>77.44+5.84</b>	<b>85.84+1.22</b>	<b>74.01+7.28</b>	<b>83.58+1.73</b>	<b>79.13+6.16</b>	<b>87.08+1.34</b>
(iii)	CSPN†	64.54	80.49	71.89	85.09	67.52	82.36	73.00	86.28
	+ ESFR	<b>71.71+7.17</b>	<b>82.22+1.73</b>	<b>78.17+6.28</b>	<b>86.38+1.29</b>	<b>74.83+7.31</b>	<b>84.17+1.81</b>	<b>79.65+6.65</b>	<b>87.57+1.29</b>
(iv)	BD-CSPN†	70.00	<b>82.36</b>	77.28	<b>86.55</b>	72.74	84.14	78.89	<b>87.72</b>
	+ ESFR	<b>73.98+3.98</b>	82.32-0.04	<b>80.13+2.85</b>	86.34-0.21	<b>76.84+4.10</b>	<b>84.36+0.22</b>	<b>81.77+2.88</b>	87.61-0.11
	+ ESFR-Semi		<b>82.89+0.53</b>		<b>86.83+0.28</b>		<b>84.97+0.83</b>		<b>88.10+0.38</b>

ing and l2-normalization to embedding network’s output for reconstruction training and baseline methods. For centering, we subtract the center of task sample embeddings from each embedding as in Lichtenstein et al. (2020) since it performs better on all baseline methods. For BD-CSPN (Liu et al., 2020), the aforementioned shifting-term defined as:

$$\Delta = \frac{1}{|S|} \sum_{x_s \in S} f(x_s) - \frac{1}{|Q|} \sum_{x_q \in Q} f(x_q) \quad (8)$$

is added for query sample before centering as in Ziko et al. (2020); Liu et al. (2020). For reconstruction training, when computing the reconstruction loss, output embeddings of the reconstruction module and the pre-trained embedding network follows the same preprocessing. For few-shot classification with new embeddings, we apply only l2-normalization since it performs the best.<sup>12</sup> We refer to Appendix A for more details on preprocessing.

### 5.3. Results

**Improvement by ESFR:** We validate the improvement of our method when used in conjunction with commonly used few-shot classification methods. We evaluate our method in the most common 5-way 1- and 5-shot settings on standard *mini-ImageNet*, *tiered-ImageNet* datasets with ResNet-18, WRN-28-10 backbone networks. The results are listed in Table 1.

In Table 1-(i, ii), we investigate our method with linear and NN classifiers. Each is a widely used baseline method in unsupervised representation learning and few-shot classification. Our method provides noticeable improvements in all settings; +5.9%~8.9% for 1-shot and +1.5%~3.1% for 5-shot settings. This roughly indicates that our embedding adaptation provides well-clustered new embeddings; hence samples can be classified by simple affine classifiers.

<sup>12</sup>Note that applying only l2-normalization for baseline methods worsen the performance.

In Table 1-(iii, iv), we compare our method with the semi-supervised learning approach, Prototype-Rectification (PR) (Liu et al., 2020). For a fair comparison, we use the same preprocessing with shifting-term (8) and cosine similarity-based nearest neighbor (CSPN) classifier as in BD-CSPN† (which includes PR); the result is described as CSPN† + ESFR. We can observe that ESFR further improves accuracies on 1-shot by +0.8%~2.1% compared to BD-CSPN† while showing comparable improvements on 5-shot. This result implies that, though unsupervised adaptation has received less attention in few-shot classification, well-designed unsupervised learning can provide comparable improvements to conventional transductive methods.

Performance can be further improved by incorporating our method into BD-CSPN, a transductive method. The result is described in Table 1-(iv). For the 5-shot setting, there was no additional gain by combining our method; however, for the 1-shot setting, improvements of +2.9%~4.1% indicates that ESFR can offer a complementary improvement to PR.

In the last row of Table 1-(iv), we also investigate ESFR-semi that uses the semi-supervised loss  $\mathcal{L}_{\text{Semi}}$  (7) for embedding adaptation. As mentioned in Section 5.2, we chose the trade-off parameter  $\lambda$  that best performs with validation tasks. For the 1-shot setting, we find that  $\lambda = 0$  performs the best that does not use label information during adaptation. For the 5-shot setting, ESFR-semi gives an additional +0.3%~0.8% gain over BD-CSPN.

**Comparison with prior work:** We compare our method with prior few-shot classification methods<sup>13</sup> on *mini-ImageNet*, *tiered-ImageNet*, and CUB datasets in Table 2. We use BD-CSPN + ESFR and BD-CSPN + ESFR-Semi as our methods.

For the 1-shot setting, our method achieves new state-of-the-

<sup>13</sup>We separately compared the work by (Boudiaf et al., 2020) in Appendix B; since they use strong prior that the number of query samples per class is equal.

Table 2. Comparison with state-of-the-art methods of 5-way 1- and 5-shot accuracy (in %) on *mini-ImageNet*, *tiered-ImageNet* and CUB. The best results are reported in **bold**.

Method	Backbone	<i>mini-ImageNet</i>		<i>tiered-ImageNet</i>		CUB	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
MAML (Finn et al., 2017)	ResNet-18	49.61	65.72	-	-	68.42	83.47
Chen (Chen et al., 2019)	ResNet-18	51.87	75.68	-	-	67.02	83.58
ProtoNet (Snell et al., 2017)	ResNet-18	54.16	73.68	-	-	72.99	86.64
TPN (Liu et al., 2019)	ResNet-12	59.46	75.65	-	-	-	-
TEAM (Qiao et al., 2019)	ResNet-18	60.07	75.90	-	-	80.16	87.17
SimpleShot (Wang et al., 2019)	ResNet-18	63.10	79.92	69.68	84.56	70.28	86.37
CTM (Li et al., 2019)	ResNet-18	64.12	78.64	68.41	84.28	-	-
FEAT (Ye et al., 2020)	ResNet-18	66.78	82.05	70.80	84.79	-	-
BD-CSPN (Liu et al., 2020)	ResNet-18	70.00	82.36	77.28	86.55	78.89	88.70
LaplacianShot (Ziko et al., 2020)	ResNet-18	72.11	82.31	78.98	86.39	80.96	88.68
BD-CSPN + ESFR (Ours)	ResNet-18	<b>73.98</b>	<b>82.32</b>	<b>80.13</b>	<b>86.34</b>	<b>82.68</b>	<b>88.65</b>
BD-CSPN + ESFR-Semi (Ours)	ResNet-18	-	<b>82.89</b>	-	<b>86.83</b>	-	<b>89.10</b>
LEO (Rusu et al., 2019)	WRN	61.76	77.59	66.33	81.44	-	-
wDAE-GNN (Gidaris & Komodakis, 2019)	WRN	62.96	78.85	68.18	83.09	-	-
FEAT (Ye et al., 2020)	WRN	65.10	81.11	70.41	84.38	-	-
Tran. Baseline (Dhillon et al., 2020)	WRN	65.73	78.40	73.34	85.50	-	-
SimpleShot (Wang et al., 2019)	WRN	65.87	82.09	70.90	85.76	-	-
SIB (Hu et al., 2020)	WRN	70.0	79.2	-	-	-	-
BD-CSPN (Liu et al., 2020)	WRN	72.74	84.14	78.89	87.72	-	-
LaplacianShot (Ziko et al., 2020)	WRN	74.86	84.13	80.18	87.56	-	-
BD-CSPN + ESFR (Ours)	WRN	<b>76.84</b>	<b>84.36</b>	<b>81.77</b>	87.61	-	-
BD-CSPN + ESFR-Semi (Ours)	WRN	-	<b>84.97</b>	-	<b>88.10</b>	-	-

Table 3. Ablation study evaluating the effects of embedding ensemble and dropout perturbation.

Method	<i>mini-ImageNet</i>		<i>tiered-ImageNet</i>		CUB	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
NN with ResNet-18	64.04	79.71	71.60	84.62	71.43	86.44
(i) w/o Dropout and Ensemble	66.87	80.59	73.39	84.60	75.46	87.02
(ii) w/o Ensemble	69.66	81.10	76.31	85.33	78.32	87.63
(iii) w/o Dropout	68.90	81.53	75.39	85.31	77.32	87.66
<b>NN + ESFR</b>	<b>70.94</b>	<b>81.61</b>	<b>77.44</b>	<b>85.84</b>	<b>79.44</b>	<b>88.02</b>

art performance on all datasets and backbone networks. To be specific, our method outperforms the previous state-of-the-art LaplacianShot (Ziko et al., 2020) by +1.2%~2.0%. We note that our implementation is based on LaplacianShot; BD-CSPN + ESFR shares many aspects including the baseline method BD-CSPN and embedding network pre-training.

For the 5-shot setting, our method shows comparable performance to the state-of-the-art. Note that in Table 1, CSPN + ESFR, which does not benefit from transductive inference, shows subequal performance with BD-CSPN + ESFR. For BD-CSPN + ESFR-Semi, our method provides further improvements of +0.3%~0.8%.

**Ablation study:** We conduct an ablation analysis on the effects of different components of the proposed method. We use a NN classifier with ResNet-18 backbone network for these experiments; our experimental results are on *mini-ImageNet*, *tiered-ImageNet*, and CUB with 5-way 1- and 5-shot tasks. Table 3 shows the influences of the embedding ensemble and dropout perturbation.

In Section 3.3, we proposed the embedding ensemble to

reduce the variance by random initialization. The empirical result in Table 3-(iii) shows that the ensemble consistently provides improvements in all settings. In Section 3.1, we expected the effectiveness of dropout perturbation since noise perturbations tend to make memorization harder. The result in Table 3-(ii) shows consistent improvements by dropout; supports our expectation. Finally, we observe that the complete version with both the embedding ensemble and dropout perturbation outperforms the other configurations.

**Comparison with prior embedding adaptation:** We compare our method with prior embedding adaptation methods that are based on affine transformations: 1) Embedding Propagation (EP) from Rodríguez et al. (2020), 2) Principal Component Analysis (PCA)- and Independent Component Analysis (ICA)<sup>14</sup>-based methods from Lichtenstein et al. (2020). For a fair comparison, we experiment with the same nearest neighbor classifier and pre-trained embeddings. We use the released official code of each method. Table 4 describes

<sup>14</sup>Lichtenstein et al. (2020) explained their method as Independent Component Analysis (ICA), but it seems the difference between their PCA- and ICA-based methods is only in the whitening; for ICA, they whiten by the projection of  $1^T z = 0$ .



the results of 1- and 5-shot settings on *mini-ImageNet* with WRN backbone.

Table 4. Comparison with prior embedding adaptation methods

	1-shot	5-shot
NN	66.73 $\pm$ 0.44	81.85 $\pm$ 0.31
NN + PCA	69.63 $\pm$ 0.50	82.28 $\pm$ 0.32
NN + EP	70.58 $\pm$ 0.47	82.73 $\pm$ 0.30
NN + ICA	72.19 $\pm$ 0.54	83.12 $\pm$ 0.32
<b>NN + ESFR (Ours)</b>	<b>74.01 <math>\pm</math> 0.51</b>	<b>83.58 <math>\pm</math> 0.31</b>

Our method outperforms both [Rodríguez et al. \(2020\)](#) and [Lichtenstein et al. \(2020\)](#) on all settings. Extracting shared or correlated features reminds us of the concept of PCA; however, compared to PCA, our method advantage from discovering non-linear patterns and correlates.

## 6. Conclusion

We propose the novel unsupervised embedding adaptation scheme based on the finding that *deep neural networks learn to generalize before memorizing*. Our method, ESFR, provides task-adapted embeddings of generalizable features by feature reconstruction training and LID-based early stopping. Experimental results show that well-designed unsupervised adaptation can consistently improve baseline methods; outperform conventional transductive methods; be further improved by joint usage with a transductive method. ESFR used in conjunction with the transductive method achieves new state-of-the-art performance on the 1-shot setting. We hope that our work will become a starting point for future unsupervised learning studies on few-shot classification.

## Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) [2021R1A2C2007518].

## References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *USENIX*, 2016.
- Amsaleg, L., Chelly, O., Furon, T., Girard, S., Houle, M. E., Kawarabayashi, K., and Nett, M. Estimating local intrinsic dimensionality. In *SIGKDD*, 2015.
- Amsaleg, L., Bailey, J., Barbe, D., Erfani, S. M., Houle, M. E., Nguyen, V., and Radovanovic, M. The vulnerability of learning to adversarial perturbation increases with intrinsic dimensionality. In *IEEE Workshop on Information Forensics and Security WIFS*, 2017.
- Ansuini, A., Laio, A., Macke, J. H., and Zoccolan, D. Intrinsic dimension of data representations in deep neural networks. In *NeurIPS*, 2019.
- Antoniou, A. and Storkey, A. J. Learning to learn by self-critique. In *NeurIPS*, 2019.
- Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A. C., Bengio, Y., and Lacoste-Julien, S. A closer look at memorization in deep networks. In *ICML*, 2017.
- Boudiaf, M., Ziko, I. M., Rony, J., Dolz, J., Piantanida, P., and Ayed, I. B. Information maximization for few-shot learning. In *NeurIPS*, 2020.
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, 2020.
- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., and Hinton, G. Big self-supervised models are strong semi-supervised learners. In *NeurIPS*, 2020.
- Chen, W., Liu, Y., Kira, Z., Wang, Y. F., and Huang, J. A closer look at few-shot classification. In *ICLR*, 2019.
- Dhillon, G. S., Chaudhari, P., Ravichandran, A., and Soatto, S. A baseline for few-shot image classification. In *ICLR*, 2020.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- Gidaris, S. and Komodakis, N. Generating classification weights with GNN denoising autoencoders for few-shot learning. In *CVPR*, 2019.
- Gidaris, S., Singh, P., and Komodakis, N. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018.
- Gidaris, S., Bursuc, A., Komodakis, N., Perez, P., and Cord, M. Boosting few-shot visual learning with self-supervision. In *ICCV*, 2019.
- Gong, S., Boddeti, V. N., and Jain, A. K. On the intrinsic dimensionality of image representations. In *CVPR*, 2019.
- Grill, J.-B., Strub, F., Alché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., et al. Bootstrap your own latent: A new approach to self-supervised learning. In *NeurIPS*, 2020.
- Hendrycks, D., Lee, K., and Mazeika, M. Using pre-training can improve model robustness and uncertainty. In *ICML*, 2019.

- Hou, R., Chang, H., Ma, B., Shan, S., and Chen, X. Cross attention network for few-shot classification. *NeurIPS*, 2019.
- Houle, M. E. Local intrinsic dimensionality I: an extreme-value-theoretic foundation for similarity applications. In *SISAP*, 2017a.
- Houle, M. E. Local intrinsic dimensionality II: multivariate analysis and distributional support. In *SISAP*, 2017b.
- Hu, S. X., Moreno, P. G., Xiao, Y., Shen, X., Obozinski, G., Lawrence, N. D., and Damianou, A. C. Empirical bayes transductive meta-learning with synthetic gradients. In *ICLR*, 2020.
- Jiang, L., Zhou, Z., Leung, T., Li, L.-J., and Fei-Fei, L. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, 2018.
- Kim, B., Choo, J., Kwon, Y.-D., Joe, S., Min, S., and Gwon, Y. Selfmatch: Combining contrastive self-supervision and consistency for semi-supervised learning. *arXiv preprint arXiv:2101.06480*, 2021.
- Kim, J., Kim, T., Kim, S., and Yoo, C. D. Edge-labeling graph neural network for few-shot learning. In *CVPR*, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Lampinen, A. K. and Ganguli, S. An analytic theory of generalization dynamics and transfer learning in deep linear networks. In *ICLR*, 2019.
- Li, H., Eigen, D., Dodge, S., Zeiler, M., and Wang, X. Finding task-relevant features for few-shot learning by category traversal. In *CVPR*, 2019.
- Lichtenstein, M., Sattigeri, P., Feris, R., Giryas, R., and Karlinsky, L. TAFSSL: task-adaptive feature sub-space learning for few-shot classification. In *ECCV*, 2020.
- Liu, J., Song, L., and Qin, Y. Prototype rectification for few-shot learning. In *ECCV*, 2020.
- Liu, Y., Lee, J., Park, M., Kim, S., Yang, E., Hwang, S. J., and Yang, Y. Learning to propagate labels: Transductive propagation network for few-shot learning. In *ICLR*, 2019.
- Ma, X., Li, B., Wang, Y., Erfani, S. M., Wijewickrema, S. N. R., Schoenebeck, G., Song, D., Houle, M. E., and Bailey, J. Characterizing adversarial subspaces using local intrinsic dimensionality. In *ICLR*, 2018a.
- Ma, X., Wang, Y., Houle, M. E., Zhou, S., Erfani, S. M., Xia, S., Wijewickrema, S. N. R., and Bailey, J. Dimensionality-driven learning with noisy labels. In *ICML*, 2018b.
- Miller, E. G., Matsakis, N. E., and Viola, P. A. Learning from one example through shared densities on transforms. In *CVPR*, 2000.
- Noroozi, M. and Favaro, P. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016.
- Oreshkin, B. N., Rodriguez, P., and Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. *NeurIPS*, 2018.
- Oymak, S., Fabian, Z., Li, M., and Soltanolkotabi, M. Generalization guarantees for neural networks via harnessing the low-rank structure of the jacobian. *arXiv preprint arXiv:1906.05392*, 2019.
- Qiao, L., Shi, Y., Li, J., Tian, Y., Huang, T., and Wang, Y. Transductive episodic-wise adaptive metric for few-shot learning. In *ICCV*, 2019.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- Ren, M., Ravi, S., Triantafillou, E., Snell, J., Swersky, K., Tenenbaum, J. B., Larochelle, H., and Zemel, R. S. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018.
- Rodríguez, P., Laradji, I., Drouin, A., and Lacoste, A. Embedding propagation: Smoother manifold for few-shot classification. In *ECCV*, 2020.
- Rusu, A. A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., and Hadsell, R. Meta-learning with latent embedding optimization. In *ICLR*, 2019.
- Snell, J., Swersky, K., and Zemel, R. S. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- Song, H., Kim, M., Park, D., and Lee, J.-G. How does early stopping help generalization against label noise? *arXiv preprint arXiv:1911.08059*, 2020.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- Stephenson, C., suchismita padhy, Ganesh, A., Hui, Y., Tang, H., and Chung, S. On the geometry of generalization and memorization in deep neural networks. In *ICLR*, 2021.
- Su, J.-C., Maji, S., and Hariharan, B. When does self-supervision improve few-shot learning? In *ECCV*, 2020.
- Sugiyama, M. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018.

- Thrun, S. and Pratt, L. *Learning to learn*. Springer Science & Business Media, 2012.
- Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J. B., and Isola, P. Rethinking few-shot image classification: a good embedding is all you need? *arXiv preprint arXiv:2003.11539*, 2020.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. *NeurIPS*, 2016.
- Wang, Y., Chao, W., Weinberger, K. Q., and van der Maaten, L. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019.
- Welinder, P., Branson, S., Mita, T., Wah, C., Schroff, F., Belongie, S., and Perona, P. Caltech-ucsd birds 200. 2010.
- Ye, H., Hu, H., Zhan, D., and Sha, F. Few-shot learning via embedding adaptation with set-to-set functions. In *CVPR*, 2020.
- Yu, X., Han, B., Yao, J., Niu, G., Tsang, I., and Sugiyama, M. How does disagreement help generalization against label corruption? In *ICML*, 2019.
- Zhai, X., Oliver, A., Kolesnikov, A., and Beyer, L. S4l: Self-supervised semi-supervised learning. In *ICCV*, 2019.
- Ziko, I. M., Dolz, J., Granger, E., and Ayed, I. B. Laplacian regularized few-shot learning. In *ICML*, 2020.