

SELF-SUPERVISED LEARNING FOR FEW-SHOT IMAGE CLASSIFICATION

Da Chen, Yuefeng Chen, Yuhong Li, Feng Mao, Yuan He[†], Hui Xue

Alibaba Group, China

{chen.cd, yuefeng.chenyf, daniel.lyh, maofeng.mf, heyuan.hy, hui.xueh@alibaba-inc.com}

ABSTRACT

Few-shot image classification aims to classify unseen classes with limited labelled samples. Recent works benefit from the meta-learning process with episodic tasks and can fast adapt to class from training to testing. Due to the limited number of samples for each task, the initial embedding network for meta-learning becomes an essential component and can largely affect the performance in practice. To this end, most of the existing methods highly rely on the efficient embedding network. Due to the limited labelled data, the scale of embedding network is constrained under a supervised learning (SL) manner which becomes a bottleneck of the few-shot learning methods. In this paper, we proposed to train a more generalized embedding network with self-supervised learning (SSL) which can provide robust representation for downstream tasks by learning from the data itself. We evaluate our work by extensive comparisons with previous baseline methods on two few-shot classification datasets (*i.e.*, MiniImageNet and CUB) and achieve better performance over baselines. Tests on four datasets in cross-domain few-shot learning classification show that the proposed method achieves state-of-the-art results and further prove the robustness of the proposed model. Our code is available at <https://github.com/phecy/SSL-FEW-SHOT>.

Index Terms— Few-shot learning, Self-supervised learning, Metric learning, Cross-domain

1. INTRODUCTION

Recent advances in deep learning techniques have made significant progress in many areas. The main reason for such success is the ability to train a deep model that can retain profound knowledge from large scale labelled dataset. This is somehow against human learning behaviour - one can easily classify objects from just a few examples with limited prior knowledge. How to computationally model such behaviour motivates the recent researches in few-shot learning, where the focus is on how to adapt the model to new data or tasks with a restricted number of instances.

One popular solution for few-shot classification is to apply a fine-tuning process on existing embedding network to adapt new classes. The main challenge here is that the fine-tuning could easily lead to overfitting, as only a few samples (1-shot or 5-shot) for each class are available. One recent proposed solution for few-shot classification is a meta-learning process, in which the dataset is divided into subsets for different meta tasks to learn how to adapt the model according to the task change. These methods highly rely on an effective pre-train embedding network.

Current methods [1, 2, 3, 4] with good performance mostly apply a ResNet12 [5] or a wide ResNet [6] as the embedding network and surpass the methods [7, 8] with deeper network. We argue that

the abandon of large network is mainly because all these methods are trained in a supervised way with limited labelled samples. In this paper, we propose to apply a much larger embedding network with self-supervised learning (SSL) to incorporate with episodic task based meta-learning. According to the evaluation presented in Section 4, the proposed method can significantly improve few-shot image classification performance over baseline methods in two common datasets. As a remark, under the the same experiment setting, the proposed method improves 1-shot and 5-shot tasks by nearly **3%** and **4%** on MiniImageNet, by nearly **9%** and **3%** on CUB. Moreover, the proposed method can gain the improvement of **15%**, **13%** and **15%**, **8%** in two tasks on MiniImageNet and CUB dataset by pretraining using more unlabeled data. We also observe that the proposed model can be robustly transferred to other datasets under a recently proposed cross-domain few-shot learning scenario [9] and achieve the state-of-the-art result (**69.69%** vs. 68.14%).

2. RELATED WORK

Few-shot learning as an active research topic has been extensively studied. In this paper, we will primarily review recent deep-learning based approaches that are more relevant to our work. A number of works aim to improve the robustness of the training process. Zhao *et al.* [10] split the features to three orthogonal parts to improve the classification performance for few-shot learning, allowing simultaneous feature selection and dense estimation. Chen *et al.* [8] propose a Self-Jig algorithm to augment the input data in few-shot learning by synthesizing new images that are either labelled or unlabeled.

A popular strategy for few-shot learning is through meta-learning (also called learning-to-learn) with multi-auxiliary tasks [11, 12, 13, 3]. The key is how to robustly accelerate the learning progress of the network without suffering from over-fitting with limited training data. Finn *et al.* propose MAML [12] to search the best initial weights through gradient descent for network training, making the fine-tuning easier. REPTILE [14] simplifies the complex computation of MAML by incorporating an L_2 loss, but still performs in high dimension space. To reduce the complexity, Rusu *et al.* propose a network called LEO [4] to learn a low dimension latent embedding of the model. Another stream of meta-learning based approaches [15, 2, 11, 13] attempt to learn a deep embedding model that can effectively project the input samples to a specific feature space. Then the samples can be classified by the nearest neighbour (NN) criterion using a pre-defined distance function. Deeper embedding backbone has also been tried. Chen *et al.* [8] propose a data augmentation method to cope with the over-fitting issue with deeper backbone embedding network. In [16], Hariharan *et al.* indicate similar observation, as claimed in this paper, that with a deeper backbone (ResNet-50) as embedding is not only costly but also not effective and proposed a novel representation regularization techniques.

[†] Corresponding Author.

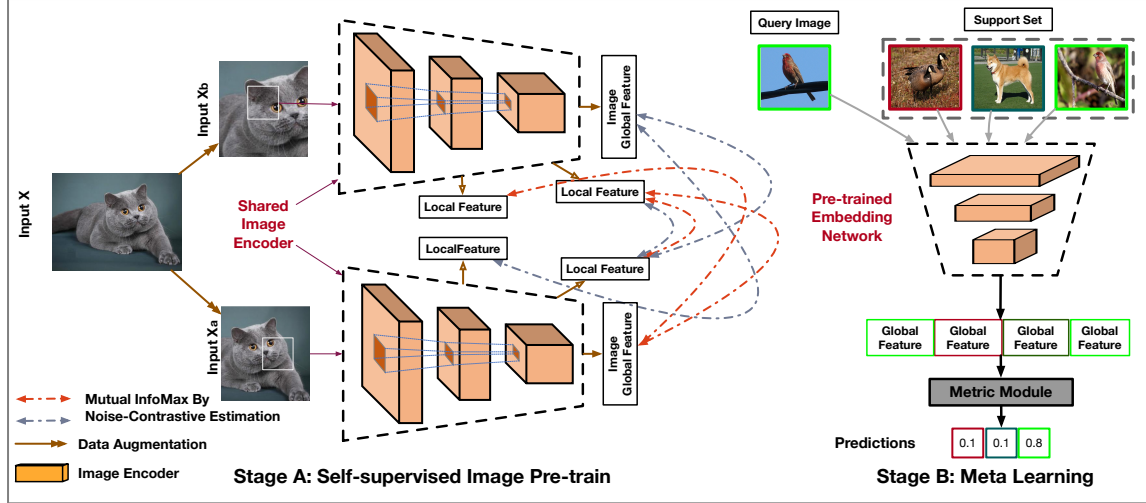


Fig. 1. The overall architecture of our approach. LEFT: Train embedding network by Self-Supervised learning. The pretext task is designed to maximize the mutual information between two views $\langle x_a, x_b \rangle$ generated from the same image x by data augmentation. Right: Meta-learning with an episodic task (3-way, 1-shot example). For each task, the training samples and query samples are encoded by the embedding network. Query sample embeddings are compared with the centroid of training sample embeddings and make a further prediction.

Self-supervised learning (SSL) aims to learn robust representations from the data itself without class labels. The main challenge here is how to design the pretext tasks that are complex enough to exploit high-level compact semantic visual representations that are useful for solving downstream tasks. This is consistent with the mission of the pre-trained embedding network in few-shot learning. The work of [17] revisited some state-of-the-art methods based on various classification based pretext tasks (e.g., Rotation, Exemplar, RelPatchLoc, Jigsaw). Recently, several methods have been proposed to combine SSL with few-shot learning. Gidaris *et al.* [18] proposed to apply a combination of supervised loss and self-supervised loss to pre-train the embedding network. A study about embedding network in few-shot learning [19] indicates that self-supervised learning based embedding network can achieve similar results as supervised learning. All these methods indicate that with a simple network such as ResNet-12, ResNet-18 can outperform ResNet-50. Different from these work, the proposed method does not rely on any supervised learning during the pre-training phase and prove that large scale embedding network can be applied in few-shot learning and achieve better results.

3. METHOD

Few-shot learning is a challenging problem as it has only limited data for training and needs to verify the performance on the data for unseen classes. An effective solution for few-shot learning classification problem is to apply a meta-learning scheme on top of a pre-trained embedding network. Most of the current methods are mainly focusing on the second stage *i.e.*, *meta-learning stage*. In this work, we follow this two stages paradigm but utilize self-supervised learning to train a large embedding network as our strong base in stage one.

3.1. Self-supervised learning stage

Our goal is to learn representations that enhance the feature's generalization. In our approach, we use Augmented Multiscale Deep InfoMax (AMDIM) [20] as our self-supervised model. The pretext

task is designed to maximize the mutual information between features extracted from multiple views of a shared context.

The mutual information (MI) measures the shared information between two random variables X and Y which is defined as the Kullback-Leibler (KL) divergence between the joint and the product of the marginals.

$$\begin{aligned}
 I(X, Y) &= D_{KL}(p(x, y) || p(x)p(y)) \\
 &= \sum \sum p(x, y) \log \frac{p(x|y)}{p(x)} \quad (1)
 \end{aligned}$$

where $P(x, y)$ is the joint distribution and $P(x)$ and $P(y)$ are the marginal distributions of X and Y . Estimating MI is challenging as we just have samples but not direct access to the underlying distribution. [21] proved that we can maximize a lower bound on mutual information by minimizing the Noise Contrastive Estimation (NCE) loss based on negative sampling.

The core is to maximize mutual information between global features and local features from two views (x_a, x_b) of the same image. Specifically, maximize mutual information between $\langle f_g(x_a), f_5(x_b) \rangle$, $\langle f_g(x_a), f_7(x_b) \rangle$ and $\langle f_5(x_a), f_5(x_b) \rangle$. Where f_g is the global feature, f_5 is encoder's 5×5 local feature map as well as f_7 as the encoder's 7×7 feature map. For example, the NCE loss between $f_g(x_a)$ and $f_5(x_b)$ is defined as below:

$$\begin{aligned}
 \mathcal{L}_{ssl}(f_g(x_a), f_5(x_b)) &= \\
 &= -\log \frac{\exp\{\phi(f_g(x_a), f_5(x_b))\}}{\sum_{\tilde{x}_b \in \mathcal{N}_x \cup x_b} \exp\{\phi(f_g(x_a), f_5(\tilde{x}_b))\}} \quad (2)
 \end{aligned}$$

\mathcal{N}_x are the negative samples of image x , ϕ is the distance metric function. At last, the overall loss between x_a and x_b is as follows:

$$\begin{aligned}
 \mathcal{L}_{ssl}(x_a, x_b) &= \mathcal{L}_{ssl}(f_g(x_a), f_5(x_b)) + \\
 &+ \mathcal{L}_{ssl}(f_g(x_a), f_7(x_b)) + \mathcal{L}_{ssl}(f_5(x_a), f_5(x_b)) \quad (3)
 \end{aligned}$$

For more details, please refer to [20].

3.2. Meta-learning stage

Given a pre-trained embedding network from stage one, meta-learning is applied to further fine-tune the model with an episodic manner. A few-shot K -way image classification task can be illustrated as a K -way C -shot problem *i.e.*, given C labelled samples for each unseen class, the model should fast adapt to them to classify novel classes. The entire training set can be presented by $D = \{(x_1, y_1), \dots, (x_N, y_N)\}$ where N is the total number of classes in D , x is the training sample with label y . For a specific K -way C -shot meta task T , $V = \{y_i | i = 1, \dots, K\}$ denotes the class labels randomly chosen from D . Training samples from these classes are randomly chosen to form a support set and a query set: **(a)** the support set for task T is denoted by S , which contains $C \times K$ samples (K -way C -shot); **(b)** the query set is Q where n is the number of samples selected for meta testing.

In this paper, during the meta-learning stage, the proposed model is trained to learn an embedding function to map all input samples from same class to a mean vector c in a description space as a class descriptor for each class [15]. For class k , it is represented by the centroid of embedding features of training samples and can be obtained as:

$$c_k = \frac{1}{|S_k|} \sum_{(x_i, y_i) \in S} f(x_i), \quad (4)$$

where $f(x_i)$ is the embedding function initialized by stage one, S_k is the training samples labelled with class k .

As a metric learning based method, we employ a distance function d and produce a distribution over all classes given a query sample q from the query set Q :

$$p(y = k | q) = \frac{\exp(-d(f(q), c_k))}{\sum_{k'} \exp(-d(f(q), c_{k'}))} \quad (5)$$

In this paper, Euclidean distance is chosen as distance function d . As shown in Eq. 5, the distribution is based on a softmax over the distance between the embedding of the samples (in the query set) and the class descriptors. The loss in the meta-learning stage can then read:

$$\mathcal{L}_{meta} = d(f(q), c_k) + \log \sum_{k'} d(f(q), c_{k'}) \quad (6)$$

In a short conclusion, the proposed method first applies an SSL way to pre-train a large scale embedding network in stage one, followed by a detailed fine-tuning in stage two with a meta-learning scheme.

4. EXPERIMENTAL RESULTS

In this section, we first introduce the dataset and training process used in our evaluation, then show quantitative comparisons against other baseline methods, finally we conduct a detailed study to validate the transferability of our approach under a cross-domain few-shot learning evaluation set-up proposed in [9].

4.1. Datasets

MiniImageNet dataset [11], is a subset of ImageNet which is a standard benchmark to evaluate the performance of few-shot learning methods. It contains 60,000 images from 100 classes, and each class has 600 images. We follow the data split strategy in [22] to sample

Baselines	Embedding Net	1-Shot 5-Way	5-Shot 5-Way
MatchingNet [11]	4 Conv	43.56 ± 0.84%	55.31 ± 0.73%
MAML [12]	4 Conv	48.70 ± 1.84%	63.11 ± 0.92%
RelationNet [13]	4 Conv	50.44 ± 0.82%	65.32 ± 0.70%
REPTILE [14]	4 Conv	49.97 ± 0.32%	65.99 ± 0.58%
ProtoNet [15]	4 Conv	49.42 ± 0.78%	68.20 ± 0.66%
Baseline* [29]	4 Conv	41.08 ± 0.70%	54.50% ± 0.66
Spot&learn [30]	4 Conv	51.03 ± 0.78%	67.96% ± 0.71
DN4 [31]	4 Conv	51.24 ± 0.74%	71.02% ± 0.64
SNAIL [32]	ResNet12	55.71 ± 0.99%	68.88 ± 0.92%
ProtoNet+ [15]	ResNet12	56.50 ± 0.40%	74.2 ± 0.20%
MTL [33]	ResNet12	61.20 ± 1.8%	75.50 ± 0.8%
DN4 [31]	ResNet12	54.37 ± 0.36%	74.44 ± 0.29%
TADAM [2]	ResNet12	58.50%	76.70%
Qiao-WRN [3]	Wide-ResNet28	59.60 ± 0.41%	73.74 ± 0.19%
LEO [4]	Wide-ResNet28	61.76 ± 0.08%	77.59 ± 0.12%
Dis. k-shot [7]	ResNet34	56.30 ± 0.40%	73.90 ± 0.30%
Self-Jig(SVM) [8]	ResNet50	58.80 ± 1.36%	76.71 ± 0.72%
FEAT [34]	ResNet50	53.8%	76.0%
Ours_Mini80_SSL	AmdimNet	43.92 ± 0.19%	67.13 ± 0.16%
Ours_Mini80_SSL ⁻	AmdimNet	46.13 ± 0.17%	70.14 ± 0.15%
Ours_Mini80_SSL	AmdimNet	64.03 ± 0.20%	81.15 ± 0.14%
Ours_Image900_SSL	AmdimNet	76.82 ± 0.19%	90.98 ± 0.10%

Table 1. Few-shot classification accuracy results on *MiniImageNet*. ‘-’ indicates result without meta-learning.

images of 64 classes for training, 16 classes for validation, 20 classes for the test.

CUB-200-2011(CUB) dataset, proposed in [23], is a dataset for fine-grained classification. It contains 200 classes of birds with 11788 images in total. For evaluation, we follow the split in [24]. 200 species of birds are randomly split to 100 classes for training, 50 classes for validation, and 50 classes for the test.

For cross-domain few-shot learning, 4 datasets are proposed to test suggested by [9], *i.e.*, 1) CropDiseases [25], a plant diseases dataset, 2) EuroSAT [26], a dataset for satellite images, 2) ISIC [27] a medical skin image dataset, 4) ChestX [28], a dataset for X-ray chest images. The similarity comparing to MiniImageNet is decreasing across these datasets.

4.2. Training Details

Several recent works show that a typical training process can include a pre-trained network [3, 4] or employ co-training [2] for feature embedding. This can significantly improve the classification accuracy. In this paper, we adopt the AMDIM [20] SSL training framework to pre-train the feature embedding network. AmdimNet(ndf=192, ndepth=8, nrkhs=1536) is used for all datasets and the embedding dimension is 1536. Adam is chosen as the optimizer with a learning rate of 0.0002. We use 128×128 as the input resolution of unlabelled images among these datasets for self-supervised training. During meta-learning stage, image size is down to 84×84 following previous methods. For MiniImageNet dataset, 3 embedding models are trained. **Mini80-SSL** is self-supervised trained from 48,000 images (80 classes training and validation) without labels. **Mini80-SL** is supervised training using same AmdimNet by cross-entropy loss with labels. **Image900-SSL** is SSL trained from all images from ImageNet1K except MiniImageNet. For CUB dataset, **CUB150-SSL** is the supervised trained model. **Image1K-SSL** is SSL trained from all images from ImageNet1K without any label. For cross-domain test, **Mini80-SSL** is applied as embedding network across all tests in four datasets during training.

Baselines	Embedding Net	1-Shot 5-Way	5-Shot 5-Way
MatchingNet [11]	4 Conv	61.16 ± 0.89	72.86 ± 0.70
MAML [12]	4 Conv	55.92 ± 0.95%	72.09 ± 0.76%
ProtoNet [15]	4 Conv	51.31 ± 0.91%	70.77 ± 0.69%
MACO [24]	4 Conv	60.76%	74.96%
RelationNet [13]	4 Conv	62.45 ± 0.98%	76.11 ± 0.69%
Baseline++ [29]	4 Conv	60.53 ± 0.83%	79.34 ± 0.61%
DN4-DA [31]	4 Conv	53.15 ± 0.84%	81.90 ± 0.60%
Ours_CUB150_SL	AmdimNet	45.10 ± 0.21%	74.59 ± 0.16%
Ours_CUB150_SSL	AmdimNet	40.83 ± 0.16%	65.27 ± 0.18%
Ours_CUB150_SSL	AmdimNet	71.85 ± 0.22%	84.29 ± 0.15%
Ours_Image1K_SSL	AmdimNet	77.09 ± 0.21%	89.18 ± 0.13%

Table 2. Few-shot classification accuracy results on CUB dataset [23]. ‘-’ indicates result without meta-learning. For each task, the best-performing method is highlighted.

4.3. Evaluation results

4.3.1. Standard few-shot learning evaluation

For MiniImageNet and CUB, we evaluate our method in two common few-shot learning tasks *i.e.*, 1-shot 5-way task and 5-shot 5-way task against baseline methods with different embedding networks including classical ones [15, 11, 12] and recently proposed methods [4, 2, 34]. For CUB dataset, we follow the work [29] to evaluate the robustness of the proposed framework with 7 other alternatives on this fine-grained dataset.

As detailed in Table 1, the proposed method outperforms all baselines in the tested tasks. In 1-shot 5-way test, our approach achieves 7.53% and 2.27% improvement over ProtoNet⁺ [15] and LEO [4] respectively. The former is an amended variant of ProtoNet using pre-trained Resnet as embedding network and has the same meta-learning stage with the proposed method. In the experience for 5-Shot 5-Way, we observe a similar improvement. Furthermore, we observe that the performance of our proposed method significantly increases when receiving more images/classes as input for pre-train. With more unlabeled samples(Image900_SSL), the model can achieve average 16% improvement over baselines while applies the same amount of labelled data.

Table 2 illustrates our experiment on CUB dataset. Our proposed method yields the highest accuracy from all trials. In the 1-shot 5-way test, we have 71.85% gaining a margin of 20.54% increment to the classic ProtoNet [15]. The improvement is more significant for the 5-shot 5-way test. Our proposed method results is 84.29% which introduces 2.39% improvement to DN4-Da [31]. Comparing to Baseline++ [29], our method shows a significant improvement, *i.e.*, 11.32% and 4.95% in both tests.

4.3.2. Cross-domain few-shot Learning

Follow the set-up in [9], the proposed method is tested on four datasets across three tasks *i.e.*, 5-way 5-shot, 5-way 20-shot 5-way 50-shot. Only Mini80_SSL embedding network(training only with MiniImageNet 80 classes) is applied in this section. We adopt the same transductive learning set-up proposed in [9] for few-shot learning. As suggested in [9], results across all tests on four datasets are averaged. As shown in Tab. 3, after averaging the test results, the proposed method has nearly 1.5% improvement over the state-of-the-art results which further prove the robustness of the proposed model. Noted that the state-of-the-art method obtain less accuracy in ISIC dataset, which needs more investigation in the future work.

Methods	ChestX			ISIC		
	5-way 5-shot	5-way 20-shot	5-way 50-shot	5-way 5-shot	5-way 20-shot	5-way 50-shot
Ours.trans	28.50 ± 0.40%	33.79 ± 0.48%	38.78 ± 0.64%	44.15 ± 0.52%	55.63 ± 0.49%	62.76 ± 0.50%
Cross [9]	26.09 ± 0.96%	31.01 ± 0.59%	36.79 ± 0.53%	49.68 ± 0.36%	61.09 ± 0.44%	67.20 ± 0.59%
	EuroSAT			CropDiseases		
	5-way 5-shot	5-way 20-shot	5-way 50-shot	5-way 5-shot	5-way 20-shot	5-way 50-shot
Ours.trans	83.44 ± 0.61%	90.43 ± 0.52%	94.71 ± 0.47%	91.79 ± 0.48%	97.38 ± 0.65%	99.50 ± 0.63%
Cross[9]	81.76 ± 0.48%	87.97 ± 0.42%	92.00 ± 0.56%	90.64 ± 0.54%	95.91 ± 0.72%	97.48 ± 0.56%

Table 3. Cross-domain few-shot learning tests on four datasets

4.4. Ablation Study

As shown in the quantitative evaluation, the proposed method can significantly improve the performance in few-shot classification task by including a large scale embedding network. One concern that may be raised is that if the gain of improvements of the proposed network is simply due to the increment of the network’s capacity. To prove the effectiveness of the proposed method, we train the embedding network with labelled data (Mini80-SL and CUB150-SL as detailed in Section 4.2). As shown in Table 1 and Table 2, it performs even worse than the methods with simple 4 Conv blocks embedding networks as such big network under supervised learning with limited data can cause overfitting problem and cannot adjust to new unseen classes during testing. However, with SSL based pre-training a more generalized embedding network can be obtained and improve the results significantly. One may also concern about the effectiveness of the meta-learning fine-tuning in the second stage. To test this, the pre-train embedding network is directly applied to the task with the nearest neighbourhood(NN) classification. As shown in the test results on both dataset, meta-learning can effectively fine-tune the embedding network and achieve remarkable improvement.

We also include more data without labels during SSL pre-training and observe a more significant improvement of the result. As shown in Table 1, the proposed method can gain the improvement of 15% and 13% in two test tasks. As detailed analyzed in [29], current few-shot learning methods can not efficiently transfer the domain of learning, *i.e.*, the training domain can not have a huge gap with the testing set. In this paper, a transferability test is also conducted by pre-training the embedding network on ImageNet and applied on CUB dataset. As shown in Table 2, the proposed method with ImageNet pre-trained embedding network can be efficiently transferred to CUB dataset and gain an improvement of 15%,8% in both test tasks. Tests on four extra datasets suggested by [9] further prove the transferability of the proposed method.

5. CONCLUSION

In this paper, we propose to utilize self-supervised learning to efficiently train a robust embedding network for few-shot image classification. The resulted embedding network is more generalized and more transferable comparing to other baselines. After fine-tuning by a meta-learning process, the performance of the proposed method can significantly outperform all baselines based on the quantitative results using two common few-shot classification datasets and a cross-domain few-shot learning set-up. The current framework can be extended in several ways in the future. For instance, one direction is to combine these two stages together and develop an end-to-end method for this task. Another direction is to investigate the effectiveness of the proposed method on another few-shot tasks such as few-shot detection, *etc.*

6. ACKNOWLEDGEMENT

This work was supported by Alibaba Group.

7. REFERENCES

- [1] Xiang Jiang, Mohammad Havaei, Farshid Varno, Gabriel Chartrand, Nicolas Chapados, and Stan Matwin, "Learning to learn with conditional class dependencies," in *ICLR*, 2019.
- [2] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste, "Tadam: Task dependent adaptive metric for improved few-shot learning," in *NeurIPS*, 2018, pp. 719–729.
- [3] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille, "Few-shot image recognition by predicting parameters from activations," in *CVPR*, 2018, pp. 7229–7238.
- [4] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell, "Meta-learning with latent embedding optimization," in *ICLR*, 2019.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *CVPR*. IEEE, 2016, pp. 770–778.
- [6] Sergey Zagoruyko and Nikos Komodakis, "Wide residual networks," *arXiv preprint arXiv:1605.07146*, 2016.
- [7] Matthias Bauer, Mateo Rojas-Carulla, Jakub Bartłomiej Swiatkowski, Bernhard Scholkopf, and Richard E Turner, "Discriminative k-shot learning using probabilistic models," *arXiv preprint arXiv:1706.00326*, 2017.
- [8] Zitian Chen, Yanwei Fu, Kaiyu Chen, and Yu-Gang Jiang, "Image block augmentation for one-shot learning.," in *AAAI*, 2019.
- [9] Yunhui Guo, Noel C Codella, Leonid Karlinsky, James V Codella, John R Smith, Kate Saenko, Tajana Rosing, and Rogério Feris, "A broader study of cross-domain few-shot learning," *ECCV*, 2020.
- [10] Bo Zhao, Xinwei Sun, Yanwei Fu, Yuan Yao, and Yizhou Wang, "Msplit lbi: Realizing feature selection and dense estimation simultaneously in few-shot and zero-shot learning," *arXiv preprint arXiv:1806.04360*, 2018.
- [11] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al., "Matching networks for one shot learning," in *NeurIPS*, 2016, pp. 3630–3638.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *ICML*. JMLR, 2017, pp. 1126–1135.
- [13] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales, "Learning to compare: Relation network for few-shot learning," in *CVPR*, 2018.
- [14] Alex Nichol and John Schulman, "Reptile: a scalable meta-learning algorithm," *arXiv preprint arXiv:1803.02999*, vol. 2, 2018.
- [15] Jake Snell, Kevin Swersky, and Richard Zemel, "Prototypical networks for few-shot learning," in *NeurIPS*, 2017.
- [16] Bharath Hariharan and Ross Girshick, "Low-shot visual recognition by shrinking and hallucinating features," in *ICCV*, 2017, pp. 3018–3027.
- [17] Alexander Kolesnikov, Xiaohua Zhai, and Lucas Beyer, "Revisiting self-supervised visual representation learning," in *CVPR*, 2019, pp. 1920–1929.
- [18] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord, "Boosting few-shot visual learning with self-supervision," in *ICCV*, 2019, pp. 8059–8068.
- [19] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola, "Rethinking few-shot image classification: a good embedding is all you need?," *arXiv preprint arXiv:2003.11539*, 2020.
- [20] Philip Bachman, R Devon Hjelm, and William Buchwalter, "Learning representations by maximizing mutual information across views," *arXiv preprint arXiv:1906.00910*, 2019.
- [21] Aaron van den Oord, Yazhe Li, and Oriol Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [22] Sachin Ravi and Hugo Larochelle, "Optimization as a model for few-shot learning," in *ICLR*, 2017.
- [23] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, "The Caltech-UCSD Birds-200-2011 Dataset," Tech. Rep. CNS-TR-2011-001, California Institute of Technology, 2011.
- [24] Nathan Hilliard, Lawrence Phillips, Scott Howland, Artëm Yankov, Courtney D Corley, and Nathan O Hodas, "Few-shot learning with metric-agnostic conditional embeddings," *arXiv preprint arXiv:1802.04376*, 2018.
- [25] Sharada P Mohanty, David P Hughes, and Marcel Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in plant science*, vol. 7, pp. 1419, 2016.
- [26] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth, "Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 7, pp. 2217–2226, 2019.
- [27] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kallou, Konstantinos Liopyris, Michael Marchetti, et al., "Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic)," *arXiv preprint arXiv:1902.03368*, 2019.
- [28] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers, "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *CVPR*, 2017, pp. 2097–2106.
- [29] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang, "A closer look at few-shot classification," in *ICLR*, 2019.
- [30] Wen-Hsuan Chu, Yu-Jhe Li, Jing-Cheng Chang, and Yu-Chiang Frank Wang, "Spot and learn: A maximum-entropy patch sampler for few-shot image classification," in *CVPR*, 2019, pp. 6251–6260.
- [31] Wenbin Li, Lei Wang, Jinglin Xu, Jing Huo, Yang Gao, and Jiebo Luo, "Revisiting local descriptor based image-to-class measure for few-shot learning," in *CVPR*, 2019, pp. 7260–7268.
- [32] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel, "A simple neural attentive meta-learner," *arXiv preprint arXiv:1707.03141*, 2017.
- [33] Qianru Sun, Yaoyao Liu, Tat-Seng Chua, and Bernt Schiele, "Meta-transfer learning for few-shot learning," in *CVPR*, 2019, pp. 403–412.
- [34] Han-Jia Ye, Hexiang Hu, De-Chuan Zhan, and Fei Sha, "Few-shot learning via embedding adaptation with set-to-set functions," in *CVPR*, 2020, pp. 8808–8817.