

Iterative label cleaning for transductive and semi-supervised few-shot learning

Michalis Lazarou¹ Tania Stathaki¹ Yannis Avrithis²

¹Imperial College London

²Inria, Univ Rennes, CNRS, IRISA

Abstract

Few-shot learning amounts to learning representations and acquiring knowledge such that novel tasks may be solved with both supervision and data being limited. Improved performance is possible by transductive inference, where the entire test set is available concurrently, and semi-supervised learning, where more unlabeled data is available.

Focusing on these two settings, we introduce a new algorithm that leverages the manifold structure of the labeled and unlabeled data distribution to predict pseudo-labels, while balancing over classes and using the loss value distribution of a limited-capacity classifier to select the cleanest labels, iteratively improving the quality of pseudo-labels. Our solution surpasses or matches the state of the art results on four benchmark datasets, namely miniImageNet, tieredImageNet, CUB and CIFAR-FS, while being robust over feature space pre-processing and the quantity of available data. The publicly available source code can be found in <https://github.com/MichalisLazarou/iLPC>

1. Introduction

Few-shot learning [61, 56] is challenging the deep learning paradigm in that, not only supervision is limited, but data is limited too. Despite the initial promise of *meta-learning* [39, 12], *transfer learning* [10, 59] is becoming increasingly successful in decoupling representation learning from learning novel tasks on limited data. *Semi-supervised learning* [30, 5] is one of the dominant ways of dealing with limited supervision and indeed, its few-shot learning counterparts [50, 66] are miniature versions where both labeled and unlabeled data are limited proportionally, while representation learning may be decoupled. These methods are closer to *transductive inference* [36, 51], which was a pillar of semi-supervised learning before deep learning [8].

Predicting *pseudo-labels* on unlabeled data [30] is one of the oldest ideas in semi-supervised learning [54]. Graph-based methods, in particular *label propagation* [68, 67], are prominent in transductive inference and translate to inductive inference in deep learning exactly by predicting pseudo-

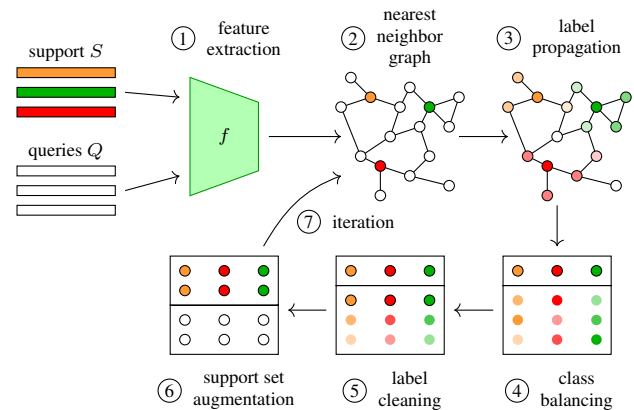


Figure 1. Overview of the proposed method. See text for details.

labels [22]. However, with the representation being fixed, the quality of pseudo-labels is critical in few-shot learning [63, 29]. At the same time, in learning with noisy labels [3, 21, 57], it is common to clean labels based on the loss value statistics of a small-capacity classifier.

In this work, we leverage these ideas to improve transductive and semi-supervised few-shot learning. As shown in Figure 1, focusing on transduction, a set of labeled support examples S and unlabeled queries Q are given, represented in a feature space by mapping f . By label propagation [67], we obtain a matrix that associates examples to classes. The submatrix corresponding to unlabeled examples, P , is normalized over examples and classes using the Sinkhorn-Knopp algorithm [24], assuming a uniform distribution over classes. We extract pseudo-labels from P , which we clean following O2U-Net [21], keeping only one example per class. Finally, inspired by [26], we move these examples from Q to S and iterate until Q is empty.

2. Related work and contributions

2.1. Few-shot learning

Meta-learning This is a popular paradigm, where the training set is partitioned in episodes resembling in struc-

ture the novel tasks [12, 25, 56, 39, 61]. *Model-based* methods rely on the properties of specific model architectures, such as recurrent and memory-augmented networks [39, 53, 40]. *Optimization-based* methods attempt to learn model parameters that are able to adapt fast in novel tasks [12, 48, 70, 49, 41, 31, 6]. *Metric-based* methods attempt to learn representations that are appropriate for comparisons [56, 25, 58, 61]. Of course, *metric learning* is a research area on its own [62, 23] and modern ideas are commonly effective in few-shot learning [34].

Predicting weights, data augmentation Also based on meta-learning, it is possible to predict new parameters or even data. For instance, it is common to learn to predict data-dependent network parameters in the last layer (classifier) [13, 45, 47] or even in intermediate convolutional layers [7]. Alternatively, one can learn to generate novel-task data in the feature space [14] or in the input (image) space [64, 2]. The quality can be improved by translating images, similar to style transfer [35]. Such learned data augmentation is complementary to other ideas.

Transfer learning More recently, it is recognized that learning a powerful representation on the entire training set is more effective than sampling few-shot training episodes that resemble novel tasks [13, 10, 59, 38]. In doing so, one may use standard loss functions [13, 10], knowledge distillation [59] or other common self-supervision and regularization methods [38]. We follow this *transfer learning* approach, which allows us to decouple representation learning from the core few-shot learning idea and provide clearer comparisons with the competition.

2.2. Using unlabeled data

Leveraging unlabelled data is of interest due to the ease of obtaining such data. Two common settings are transductive inference and semi-supervised learning.

Transductive inference In this setting, all novel-class unlabeled query examples are assumed available at the same time at inference [36, 33, 46, 18, 19, 51]. These examples give additional information on the distribution of novel classes on top of labeled support examples.

Common transductive inference solutions are adapted for few-shot classification, notably *label propagation* [36] and *embedding propagation* [51], which smooths embeddings as in image segmentation [4]. Both operations are also used at representation learning, as in meta-learning. Using dimensionality reduction, TAFSSL [33] learns a task-specific feature subspace that is highly discriminant for novel tasks. *Meta-confidence transduction* (MCT) [29] meta-learns a data-dependent scaling function to normalize every example and iteratively updates class centers. PT+MAP [19] uses a similar iterative process but also balances over classes. *Cross-attention* [16], apart from aligning feature maps by

correlation, leverages query examples by iteratively making predictions and using the most confident ones to update the class representation.

Semi-supervised learning In this case, labeled novel-class support examples and additional unlabeled data are given. A classifier may be learned on both to make predictions on novel-class queries [50, 63, 66].

One of the first contributions uses unlabelled examples to adapt prototypical networks [56], while discriminating from distractor classes [50]. Common semi-supervised solutions are also adapted to few-shot classification, for instance *learning to self-train* [32], which adapts *pseudo-label* [30] and TransMatch [66], which is an adaptation of MixMatch [5]. *Instance credibility inference* [63] predicts pseudo-labels iteratively, using a linear classifier to select the most likely to be correct and then augmenting the support set. *Adaptive subspaces* [55] are learned from labeled and unlabeled data, yielding a discriminative subspace classifier that maximizes the margin between subspaces.

2.3. Contributions

In this work, focusing on the transfer learning paradigm to learn novel tasks given a fixed representation [59, 38], we make the following contributions:

1. We combine the power of predicting pseudo-labels in semi-supervised learning [30, 32] with label cleaning in learning from noisy labels [21].
2. According to manifold assumption, we use label propagation [67, 36] to infer pseudo-labels, while balancing over classes [24, 19] and iteratively re-use pseudo-labels in the propagation process [26, 63].
3. We achieve new state of the art in both transductive and semi-supervised few-shot learning.

3. Method

3.1. Problem formulation

At *representation learning*, we assume access to a labeled dataset D_{base} with each example having a label in one of the classes in C_{base} . This dataset is used to learn a mapping $f : \mathcal{X} \rightarrow \mathbb{R}^d$ from an input space \mathcal{X} to a d -dimensional *feature* or *embedding* space.

The knowledge acquired at representation learning is used to solve *novel tasks*, assuming access to a dataset D_{novel} with each example being associated with one of the classes C_{novel} , where C_{novel} is disjoint from C_{base} . Examples in D_{novel} may be labeled or not.

In *few-shot classification* [61], a novel task is defined by sampling a *support set* S from D_{novel} , consisting of N classes with K labeled examples per class, for a total of $L := NK$ examples. Given the mapping f and the support set S , the problem is to learn an N -way classifier that makes

predictions on unlabeled queries also sampled from D_{novel} . Queries are treated independently of each other. This is referred to as *inductive inference*.

In *transductive inference*, a *query set* Q consisting of M unlabeled examples is also sampled from D_{novel} . Given the mapping f , S and Q , the problem is to make predictions on Q , without necessarily learning a classifier. In doing so, one may leverage the distribution of examples in Q , which is important because M is assumed greater than L .

In *semi-supervised* few-shot classification, an unlabelled set U of M unlabeled examples is also sampled from D_{novel} . Given f , S and U , the problem is to learn to make predictions on new queries from D_{novel} , as in inductive inference. Again, $M > L$ and we may leverage the distribution of U .

In this work, we focus on transductive inference and semi-supervised classification, given f . The performance of f on inductive inference is one of our baselines. We develop our solution for transductive inference. In the semi-supervised case, we follow the same solution with Q replaced by U . Using the predictions on U , we then proceed as in the inductive case, with S replaced by $S \cup U$.

3.2. Nearest-neighbor graph

We are given the mapping f , the labeled support set $S := \{(x_i, y_i)\}_{i=1}^L$ and the query set $Q := \{x_{L+i}\}_{i=1}^M$, where $y_i \in [N] := \{1, \dots, N\}$. We embed all examples from S and Q into $V = \{\mathbf{v}_1, \dots, \mathbf{v}_T\} \subset \mathbb{R}^d$ and ℓ_1 -normalize them, where $T := L + M$ and $\mathbf{v}_i := f(x_i)$ for $i \in [T]$. Following [22], we construct a k -nearest neighbour graph of the features in V , represented by a sparse $T \times T$ nonnegative *affinity matrix* A , with

$$A_{ij} := \begin{cases} [\mathbf{v}_i^\top \mathbf{v}_j]^\gamma, & \text{if } i \neq j \wedge \mathbf{v}_i \in \text{NN}_k(\mathbf{v}_j) \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

for $i \in [T]$, $j \in [N]$, where $\text{NN}_k(\mathbf{v})$ are the k -nearest neighbors of \mathbf{v} in V and $\gamma > 1$ is a hyperparameter. Finally, we obtain the symmetric $T \times T$ *adjacency matrix* $W := \frac{1}{2}(A + A^\top)$ and we symmetrically normalize it as

$$\mathcal{W} := D^{-1/2} W D^{-1/2}, \quad (2)$$

where $D = \text{diag}(W \mathbf{1}_T)$ is the $T \times T$ degree matrix of W .

3.3. Label propagation

Following [67], we define the $T \times N$ *label matrix* Y as

$$Y_{ij} := \begin{cases} 1, & \text{if } i \leq L \wedge y_i = j \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

for $i \in [T]$, $j \in [N]$. Matrix Y has one column per class and one row per example, which is an one-hot encoded label

for S and a zero vector for Q . Label propagation amounts to solving N linear systems

$$Z := (I - \alpha \mathcal{W})^{-1} Y, \quad (4)$$

where $\alpha \in [0, 1)$ is a hyperparameter. The resulting $T \times N$ matrix Z can be used to make predictions by taking the maximum element per row [67]. However, before making predictions, we balance over classes.

3.4. Class balancing

We focus on the $M \times N$ submatrix

$$P := Z_{L+1:T,:}, \quad (5)$$

(the last M rows) of Z that corresponds to unlabeled queries. We first perform an element-wise *power transform*

$$P_{ij} \leftarrow P_{ij}^\tau \quad (6)$$

for $i \in [M]$, $j \in [N]$, where $\tau > 1$, encouraging hard predictions. Parameter τ is analogous to the *scale* (or *inverse temperature*) of logits in softmax-based classifiers [13, 45, 42], only here the elements of P are proportional to class probabilities rather than logits.

Inspired by [19], we normalize P to a given row-wise sum $\mathbf{p} \in \mathbb{R}^M$ and column-wise sum $\mathbf{q} \in \mathbb{R}^N$. Each element $p_i \in [0, 1]$ of \mathbf{p} represents a confidence of example x_{L+i} for $i \in [M]$; it can be a function of the i -th row of P or set to 1. Each element $q_j \geq 0$ of \mathbf{q} represents a weight of class j for $j \in [N]$. In the absence of such information, we set

$$\mathbf{q} := \frac{1}{N} (\mathbf{p}^\top \mathbf{1}_M) \mathbf{1}_N, \quad (7)$$

assuming a uniform distribution of queries over classes.

The normalization itself is a projection of P onto the set $\mathbb{S}(\mathbf{p}, \mathbf{q})$ of nonnegative $M \times N$ matrices having row-wise sum \mathbf{p} and column-wise sum \mathbf{q} ,

$$\mathbb{S}(\mathbf{p}, \mathbf{q}) := \{X \in \mathbb{R}^{M \times N} : X \mathbf{1}_N = \mathbf{p}, X^\top \mathbf{1}_M = \mathbf{q}\}. \quad (8)$$

We use the *Sinkhorn-Knopp* algorithm [24] for this projection, which alternates between rescaling the rows of P to sum to \mathbf{p} and its columns to sum to \mathbf{q} ,

$$P \leftarrow \text{diag}(\mathbf{p}) \text{diag}(P \mathbf{1}_N)^{-1} P \quad (9)$$

$$P \leftarrow P \text{diag}(P^\top \mathbf{1}_M)^{-1} \text{diag}(\mathbf{q}), \quad (10)$$

until convergence. Finally, for each query x_{L+i} , $i \in [M]$, we predict the *pseudo-label*

$$\hat{y}_{L+i} := \arg \max_{j \in [N]} P_{ij} \quad (11)$$

that corresponds to the maximum element of the i -th row of the resulting matrix P , for $i \in [M]$.

3.5. Label cleaning

The predicted pseudo-labels are not necessarily correct, yet a classifier can be robust to such noise. This is the case when enough data is available to adapt the representation [30, 22], such that the quality of pseudo-labels improves with training. Since data is limited here, we would like to select pseudo-labeled queries in Q that are most likely to be correct, treat them as truly labeled and add them to the support set S . Iterating this process is an alternative way of improving the quality of pseudo-labels.

We interpret this problem as *learning with noisy labels*, leveraging recent advances in label cleaning [3, 21, 57]. Assuming that the classifier does not overfit the data, e.g. with small capacity, high learning rate or few iterations, the principle is that examples with clean labels exhibit less loss than examples with noisy labels.

In particular, given the labeled support set $S := \{(x_i, y_i)\}_{i=1}^L$ and the pseudo-labeled query set $\hat{Q} := \{(x_{L+i}, \hat{y}_{L+i})\}_{i=1}^M$, we train an N -way classifier g using a weighted cross-entropy loss

$$\ell := - \sum_{i=1}^L \log g(x_i)_{y_i} - \sum_{i=1}^M p_i \log g(x_{L+i})_{\hat{y}_{L+i}}, \quad (12)$$

where p_i is the confidence weight of example x_{L+i} . Here, the classifier g is assumed to yield a vector of probabilities over classes using softmax and $g(x)_y$ refers to element $y \in [N]$ of $g(x)$. In practice, it is obtained by a linear classifier on top of embedding f , optionally allowing the adaptation of the last layers of the network implementing f .

The loss term $\ell_i := -p_i \log g(x_{L+i})_{\hat{y}_{L+i}}$, corresponding to the pseudo-labeled query x_{L+i} , is used for selection. Following O2U-Net [21], we use large learning rate and collect the average loss $\bar{\ell}_i$ over all epochs, for $i \in [M]$. In learning with noisy labels, it is common to detect noisy labels based on statistics of this loss for clean and noisy labels [3, 57]. However, this does not work well with predicted pseudo-labels [1], hence we select queries having the least average loss [21, 1]. Since we iterate the process, we take the extreme case of selecting one query example per class:

$$\mathcal{I} := \left\{ \arg \min_{\hat{y}_{L+i}=j} \bar{\ell}_i : j \in [N] \right\}. \quad (13)$$

Finally, we augment the support set S with the selected queries and their pseudo-labels, while at the same time removing the selected queries from Q :

$$S \leftarrow S \cup \{(x_{L+i}, \hat{y}_{L+i})\}_{i \in \mathcal{I}} \quad (14)$$

$$Q \leftarrow Q \setminus \{x_{L+i}\}_{i \in \mathcal{I}} \quad (15)$$

3.6. Iterative inference

Although label propagation and class balancing make predictions on the entire unlabeled query set Q , we apply

cleaning to only keep one pseudo-labeled query per class, which we move from Q to the support set S . We iterate the entire process, selecting one pseudo-labeled query per class at a time, until Q is empty and S is augmented with all pseudo-labeled queries. Assuming that the selections are correct, the idea is that treating them as truly labeled in S improves the quality of the pseudo-labels.

Algorithm 1: Iterative label propagation and cleaning (iLPC).

```

input : embedding  $f$ 
input : labeled support set  $S$  with  $|S| = L$ 
input : unlabeled query set  $Q$  with  $|Q| = M$ 
output : augmented support set  $S$  with  $|S| = L + M$ 

1 repeat
2    $\mathcal{W} \leftarrow \text{GRAPH}(f, S, Q; \gamma, k)$   $\triangleright$  adjacency matrix (1),(2)
3    $Y \leftarrow \text{LABEL}(S)$   $\triangleright$  label matrix (3)
4    $Z \leftarrow \text{LP}(\mathcal{W}, Y; \alpha)$   $\triangleright$  label propagation (4)
5    $P \leftarrow Z_{L+1:L+M,:}$   $\triangleright$  unlabeled submatrix (5)
6    $P \leftarrow \text{POWER}(P; \tau)$   $\triangleright$  power transform (6)
7    $(\mathbf{p}, \mathbf{q}) \leftarrow \text{BALANCE}(P)$   $\triangleright$  class balance (7)
8    $P \leftarrow \text{SINKHORN}(P; \mathbf{p}, \mathbf{q})$   $\triangleright$  Sinkhorn-Knopp (9),(10)
9    $\hat{Y} \leftarrow \text{PREDICT}(P)$   $\triangleright$  pseudo-labels (11)
10   $\mathcal{I} \leftarrow \text{CLEAN}(f, S, Q, \hat{Y}, \mathbf{p})$   $\triangleright$  label cleaning (12),(13)
11   $(S, Q) \leftarrow \text{AUGMENT}(S, Q, \mathcal{I})$   $\triangleright$  augment support (14),(15)
12 until  $Q = \emptyset$   $\triangleright$  all queries are predicted
```

Algorithm 1 summarizes this process, called *iterative label propagation and cleaning* (iLPC). Given S , Q and the embedding f , we construct the nearest neighbor graph represented by the normalized adjacency matrix \mathcal{W} (1),(2) and we perform label propagation on the current label matrix Y (4). Focusing on the unlabeled submatrix P of the resulting matrix Z , we perform power transform (6) and row/column normalization to balance over classes (9),(10). We predict pseudo-labels \hat{Y} from the normalized P (11), which we use along with S and Q to train a linear classifier on top of f with cross entropy loss (12) and a cyclical learning rate schedule [21]. We select one query per class with the least average loss over all epochs (13), which we move from Q to S as labeled (14),(15). With Q , S redefined, we repeat the process until Q is empty.

At termination, all data is labeled in S . The predicted labels over the original queries are the output in the case of transductive inference. In semi-supervised classification, we use S to learn a new classifier and make predictions on new queries, as in inductive inference.

4. Experiments

4.1. Setup

Datasets We use four common few-shot classification benchmark datasets, *miniImageNet* [61, 49], *tieredImageNet* [9], CUB [10, 15] and CIFAR-FS [10, 27]. More details are provided in the supplementary material.

INFERENCE	COMPONENTS				RESNET-12A		WRN-28-10		
	LP	Balance	iLC	iProb	Class	1-shot	5-shot	1-shot	5-shot
Inductive						56.30 \pm 0.62	75.59 \pm 0.47	68.17 \pm 0.60	84.33 \pm 0.43
Transductive	✓					61.09 \pm 0.70	75.32 \pm 0.50	74.24 \pm 0.68	84.09 \pm 0.42
Transductive	✓	✓				65.04 \pm 0.75	76.82 \pm 0.50	79.42 \pm 0.69	85.34 \pm 0.43
Transductive	✓		✓			65.57 \pm 0.89	78.03 \pm 0.54	78.29 \pm 0.76	88.02 \pm 0.41
Transductive†	✓	✓	✓			69.79 \pm 0.99	79.82 \pm 0.55	83.05 \pm 0.79	88.82 \pm 0.42
Transductive	✓	✓		✓		58.27 \pm 0.91	74.11 \pm 0.56	80.75 \pm 0.76	87.62 \pm 0.44
Transductive		✓	✓		✓	68.79 \pm 0.96	79.93 \pm 0.56	82.04 \pm 0.78	88.89 \pm 0.41

Table 1. *Ablation study of algorithmic components* of our method iLPC on *miniImageNet*. Inductive: baseline using only support examples. LP: label propagation. Balance: class balancing (7). iLC: iterative label cleaning, without which we just output predictions (11). iProb: iterative selection of top examples per class directly as column-wise maxima of P (5) instead of iLC. Class: linear classifier used for prediction instead of LP, as in [63], with balancing still applied on output probabilities. †: default setting of iLPC.

METHOD	NETWORK	<i>miniImageNet</i>		<i>tieredImageNet</i>		CIFAR-FS		CUB	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
LR+ICI [63]	ResNet-12A	66.80	79.26	80.79	87.92	73.97	84.13	88.06	92.53
LR+ICI [63]*	ResNet-12A	66.85 \pm 0.92	78.89 \pm 0.55	82.40 \pm 0.84	88.80 \pm 0.50	75.36 \pm 0.97	84.57 \pm 0.57	86.53 \pm 0.79	92.11 \pm 0.35
iLPC (ours)	ResNet-12A	69.79 \pm 0.99	79.82 \pm 0.55	83.49 \pm 0.88	89.48 \pm 0.47	77.14 \pm 0.95	85.23 \pm 0.55	89.00 \pm 0.70	92.74 \pm 0.35
PT+MAP [19]	WRN-28-10	82.92 \pm 0.26	88.82 \pm 0.13	-	-	87.69 \pm 0.23	90.68 \pm 0.15	91.55 \pm 0.19	93.99 \pm 0.10
PT+MAP [19]*	WRN-28-10	82.88 \pm 0.73	88.78 \pm 0.40	88.15 \pm 0.71	92.32 \pm 0.40	86.91 \pm 0.72	90.50 \pm 0.49	91.37 \pm 0.61	93.93 \pm 0.32
LR+ICI [63]*	WRN-28-10	80.61 \pm 0.80	87.93 \pm 0.44	86.79 \pm 0.76	91.73 \pm 0.40	84.88 \pm 0.79	89.75 \pm 0.48	90.18 \pm 0.65	93.35 \pm 0.30
iLPC (ours)	WRN-28-10	83.05 \pm 0.79	88.82 \pm 0.42	88.50 \pm 0.75	92.46 \pm 0.42	86.51 \pm 0.75	90.60 \pm 0.48	91.03 \pm 0.63	94.11 \pm 0.30

Table 2. *Transductive inference*, comparison with LR+ICI [63] and PT+MAP [19]. *: our reproduction with official code on our datasets.



Figure 2. (a,b) Distributions of loss values (12) for correctly and incorrectly labeled examples, normalized independently. Uniform label noise: (a) 20%, (b) 40%. Pseudo-labels predicted by (11) for two different 1-shot transductive *miniImageNet* tasks (c,d).

Tasks We consider N -way, K -shot classification tasks with $N = 5$ randomly sampled novel classes and $K \in \{1, 5\}$ randomly selected examples per class for support set S , that is, $L = 5K$ examples in total. For the query set Q , we




Figure 3. *1-shot transductive inference* on *miniImageNet*, ablation over LR+ICI [63] pre-processing: dimension reduction by PCA.

randomly sample 15 additional examples per class, that is, $M = 75$ examples in total, which is the most common choice in the literature [36, 32, 66].

In the semi-supervised setting, the unlabeled set U contains an additional number of randomly sampled examples per novel class. This number depends on K . We use two settings, namely 30/50 and 100/100, where the first number (30 or 100) refers to 1-shot and the second (50 or 100) to 5-shot. Again, these are the two most common choices in semi-supervised few shot learning [32, 63, 29, 50, 66].

Unless otherwise stated, we use 1000 tasks and report mean accuracy and 95% confidence interval on the test set.

Competitors As discussed in the supplementary material, there are several flaws in experimental evaluation in the liter-

METHOD	NETWORK	miniIMAGENET		tieredIMAGENET		CIFAR-FS	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
LR+ICI [63]*	WRN-28-10	82.38 \pm 0.86	88.78 \pm 0.39	88.59 \pm 0.74	92.11 \pm 0.39	86.39 \pm 0.79	90.02 \pm 0.49
PT+MAP [19]*	WRN-28-10	83.79 \pm 0.71	88.94 \pm 0.33	88.87 \pm 0.64	92.01 \pm 0.36	87.63 \pm 0.66	90.15 \pm 0.46
iLPC (ours)	WRN-28-10	85.98 \pm 0.74	90.54 \pm 0.31	90.02 \pm 0.70	92.94 \pm 0.37	88.54 \pm 0.68	90.92 \pm 0.46

Table 3. *Transductive inference, 50 queries per class.* *: our reproduction with official code on our datasets.

METHOD	mIN	tIN	CIFAR-FS	CUB
LR+ICI [63]*	88.69 \pm 0.38	91.88 \pm 0.41	90.23 \pm 0.45	93.66 \pm 0.28
PT+MAP [19]*	89.97 \pm 0.34	93.33 \pm 0.34	91.30 \pm 0.45	94.24 \pm 0.28
iLPC (ours)	90.51 \pm 0.35	93.61 \pm 0.38	91.59 \pm 0.44	94.75 \pm 0.26

Table 4. *Transductive 10-shot inference* using WRN-28-10. mIN: miniImageNet. tIN: tieredImageNet. *: our reproduction with official code on our datasets.

METHOD	PRE	miniImageNet		tieredImageNet	
		1-shot	5-shot	1-shot	5-shot
PT+MAP [19]*		48.57 \pm 0.81	75.67 \pm 0.82	49.67 \pm 0.77	88.32 \pm 0.50
iLPC (ours)		78.89 \pm 0.90	86.80 \pm 0.46	86.52 \pm 0.47	91.07 \pm 0.47
PT+MAP [19]*	✓	82.88 \pm 0.73	88.78 \pm 0.40	88.15 \pm 0.71	92.32 \pm 0.40
iLPC (ours)	✓	83.05 \pm 0.79	88.82 \pm 0.42	88.50 \pm 0.75	92.46 \pm 0.42

Table 5. *Transductive inference, ablation over PT+MAP [19] pre-processing.* PRE: power transform, normalization, centering. *: our reproduction with official code on our datasets.

METHOD	NETOWRK	miniIMAGENET	
		1-shot	5-shot
MCT (instance,flip) [29]	ResNet-12B	78.55 \pm 0.86	86.03 \pm 0.42
MCT (no scale) [29]*	ResNet-12B	67.26 \pm 0.60	81.90 \pm 0.43
iLPC (ours)	ResNet-12B	75.58 \pm 1.16	81.58 \pm 0.50
iLPC (ours)	ResNet-12A	69.79 \pm 0.99	79.82 \pm 0.55

Table 6. *Transductive inference, comparison with MCT [29] using ResNet-12B.* *: our reproduction with official code on our datasets, without augmentation and without scaling.

ature, like the use of different networks, training, versions of datasets, dimensionality and feature pre-processing. Fair comparison is impossible, unless one uses public code to reproduce results under exactly the same setup.

In this work, we do provide completely fair comparisons with such reproduced results of three state-of-the-art methods: LR+ICI [63], PT+MAP [19] and MCT [29]. Only [63] is published, while the other two are pre-prints.

Networks We use publicly available pre-trained backbone convolutional neural networks that are trained on the base-class training set. We experiment with two popular networks, namely, the residual network ResNet-12 [42] and the wide residual network WRN-28-10 [52].

In particular, to compare with [63], we use pre-trained weights of the ResNet-12 provided by [63], which we call ResNet-12A, as well as official public code¹ for testing. To compare with [19], we use pre-trained weights of a WRN-28-10 provided by [38]², which are the same used by [19], as well as official public code³ for testing. To compare with [29], we use official public code⁴ to train from scratch another version of ResNet-12 used by [29], which we call ResNet-12B, as well as the same code for testing.

Feature pre-processing Each method uses its own feature pre-processing. LR+ICI [63] uses ℓ_2 -normalization and PCA to reduce ResNet-12A to 5 dimensions. PT+MAP [19] uses element-wise power transform, ℓ_2 -normalization and centering of WRN-28-10 features. MCT [29] uses flattening of the output tensor of ResNet-12B rather than spatial pooling. By default, we use the same choices as [19, 29] for WRN-28-10 and ResNet-12B. For ResNet-12A however, we use ℓ_2 -normalization only on transductive inference and we do not use any dimensionality reduction.

Implementation details We use PyTorch [43] and scikit-learn [44]. Label cleaning is based on a linear classifier on top of f , initialized by imprinting the average of support features per class and then trained using (12). We use SGD with momentum 0.9 and weight decay 0.0005. We use a learning rate of η for 1000 iterations. For inductive (resp. semi-supervised) learning, we use logistic regression on support (resp. also pseudo-labeled) examples, learned using scikit-learn [63]. The row-wise sum \mathbf{p} (9) is fixed to 1. The supplementary material includes more choices. It also includes inference time comparisons.

4.2. Ablation study

Hyperparameters Our hyperparameters include γ and k used in the nearest neighbor graph (1), α in label propagation (4), τ in balancing (7) and the learning rate η of label cleaning. We optimize them on the validation set of every dataset. Common choices for k and α are in [15, 20] and in [0.5, 0.8], respectively. We set $\gamma = 3$, $\tau = 3$ and $\eta = 0.1$. More details and precise choices per dataset are given in the supplementary material.

¹<https://github.com/Yikai-Wang/ICI-FSL>

²https://github.com/nupurkmr9/S2M2_fewshot

³<https://github.com/yhu01/PT-MAP>

⁴<https://github.com/seongmin-kye/MCT>

BALANCING	NETWORK	<i>mini</i> IMAGENET		<i>tiered</i> IMAGENET		CIFAR-FS		CUB	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
None	WRN-28-10	78.06 \pm 0.82	87.80 \pm 0.42	86.04 \pm 0.73	90.74 \pm 0.46	85.32 \pm 0.76	89.64 \pm 0.48	89.67 \pm 0.64	92.98 \pm 0.31
Uni	WRN-28-10	77.50 \pm 0.78	83.68 \pm 0.39	83.02 \pm 0.67	86.17 \pm 0.40	81.47 \pm 0.69	84.83 \pm 0.45	85.22 \pm 0.57	87.99 \pm 0.28
True	WRN-28-10	82.68 \pm 0.82	89.07 \pm 0.41	89.17 \pm 0.70	92.67 \pm 0.44	87.32 \pm 0.74	90.92 \pm 0.48	91.24 \pm 0.60	94.14 \pm 0.30

Table 7. *Unbalanced transductive inference* with our iLPC. Number of queries per class drawn uniformly at random from 10, ..., 20. None: no balancing. Uni: Uniform distribution. True: True distribution.

Algorithmic components Table 1 ablates our method in the presence or not of individual components, as well as using alternative components. The use of queries with label propagation gives a gain of transductive over inductive inference, up to 6% in 1-shot, while being on par with the linear classifier in 5-shot. In 1-shot, balancing and iterative label cleaning bring another gain of 4-5% each independently, while the combination of the two brings 8-9%. The performance of iterative label cleaning is further justified by its superior of performance when compared to selecting examples based on P instead.

4.3. Label cleaning: loss distribution

To illustrate our label cleaning, we conduct two experiments, showing the distribution of the loss value (12). In the first, shown in Figure 2(a,b), we inject label noise uniformly at random to the 20% (a) and 40% (b) of 500 labeled examples. The correctly and incorrectly labeled examples have very different loss distributions. Importantly, while previous work on noisy labels [3, 21, 57] attempts to detect clean examples by an optimal threshold on the loss value, we only need few clean examples per iteration. Examples with minimal loss value are clean.

The second experiment is on two novel 1-shot transductive tasks, shown in Figure 2(c,d). We use 50 unlabeled queries per class and we predict pseudo-labels according to (11). Label cleaning is more challenging here because the two distributions are more overlapping. This is natural because predictions are more informed than uniform, even if incorrect. Still, a large proportion of clean examples have a smaller loss value than the minimal value of noisy ones.

4.4. Effectiveness of class balancing

To show the effectiveness of our class balancing module, we carry out experiments in a novel setting for unbalanced few-shot transductive inference. In this setting, the number of queries per class for every few-shot task is drawn uniformly at random from $\{10 \dots 20\}$. We use no balancing, or we use balancing with uniform class distribution (7), or, assuming the prior class distribution $\mathbf{u} \in \mathbb{R}^N$ is known, we replace $\frac{1}{N}$ in (7) by \mathbf{u} . As shown in Table 7, balancing improves accuracy by a large margin, but only if the prior class distribution is known, otherwise it is harmful.

4.5. Transductive inference

Table 2 compares our iLPC with LR+ICI [63] and PT+MAP [19] under the standard setting of 15 unlabelled queries per class. The truly fair comparison is with our reproductions, indicated by *. Apart from the default networks, we also use WRN-28-10 with LR+ICI [63], since it is more powerful. Our iLPC is on par with PT+MAP [19] under this setting and superior to LR+ICI [63] by up to 3% on *miniImageNet* 1-shot.

We also experiment with 50 unlabeled queries per class, or $M = 250$ in total. As shown in Table 3, the gain over PT+MAP [19] increases significantly, up to 2% on *miniImageNet* 1-shot. This can be attributed to the fact that PT+MAP [19] operates on Euclidean space, while we capture the manifold structure, which manifests itself in the presence of more data. A 10-shot experiment, is shown in Table 4. The gain is around 0.5%.

Table 5 shows that PT+MAP [19] is very sensitive to feature pre-processing, losing up to 40% without it, while our iLPC more robust, losing only up to 5%. Similarly, Figure 3 shows that LR+ICI [63] is sensitive to dimension reduction, working best at only 5 dimensions. By contrast, our iLPC is very stable and only fails at 2 dimensions.

Table 6 compares our iLPC with MCT [29]. We reproduce MCT results by training from scratch ResNet-12B using the official code and we test both methods without data augmentation (horizontal flipping) and without meta-learned scaling function. The objective is to compare the two transductive methods under the same backbone network and the same training process, which is clearly superior to ResNet-12A. Under these settings, MCT is slightly better in 5-shot but iLPC outperforms it by a large margin in 1-shot.

4.6. Semi-supervised learning

As shown in Table 8, iLPC is superior to LR+ICI [63] in all settings by an even larger margin than in transductive inference, e.g. by nearly 3.5% in *miniImageNet* 1-shot. This can be attributed to capturing the manifold structure of the data, since there is more unlabeled data in this case. Because PT+MAP [19] does not experiment with semi-supervised learning, we adapt it in the same way as ours, using the default WRN-28-10, outperforming it in most experiments.

METHOD	NETWORK	SPLIT	miniIMAGENET		tieredIMAGENET		CIFAR-FS		CUB	
			1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
LR+ICI [63]	ResNet-12A	30/50	69.66	80.11	84.01	89.00	76.51	84.32	89.58	92.48
LR+ICI [63]*	ResNet-12A	30/50	67.57 \pm 0.97	79.07 \pm 0.56	83.32 \pm 0.87	89.06 \pm 0.51	75.99 \pm 0.98	84.01 \pm 0.62	88.50 \pm 0.71	-
iLPC (ours)	ResNet-12A	30/50	70.99 \pm 0.91	81.06 \pm 0.49	85.04 \pm 0.79	89.63 \pm 0.47	78.57 \pm 0.80	85.84 \pm 0.56	90.11 \pm 0.64	-
LR+ICI [63]*	WRN-28-10	30/50	81.31 \pm 0.84	88.53 \pm 0.43	88.48 \pm 0.67	92.03 \pm 0.43	86.03 \pm 0.77	89.57 \pm 0.53	90.82 \pm 0.59	-
PT+MAP [19]†	WRN-28-10	30/50	83.14 \pm 0.72	88.95 \pm 0.38	89.16 \pm 0.61	92.30 \pm 0.39	87.05 \pm 0.69	89.98 \pm 0.49	91.52 \pm 0.53	-
iLPC (ours)	WRN-28-10	30/50	83.58 \pm 0.79	89.68 \pm 0.37	89.35 \pm 0.68	92.61 \pm 0.39	87.03 \pm 0.72	90.34 \pm 0.50	91.69 \pm 0.55	-

Table 8. *Semi-supervised few-shot learning*, comparison with [63, 19]. *: our reproduction with official code on our datasets. †: our adaptation to semi-supervised, based on official code. CUB 5-shot omitted: no class has the required 70 examples. ‡: with centering.

METHOD	NETWORK	miniIMAGENET		tieredIMAGENET		CIFAR-FS		CUB	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
LR+ICI [63]*	ResNet-12A	66.85 \pm 0.92	78.89 \pm 0.55	82.40 \pm 0.84	88.80 \pm 0.50	75.36 \pm 0.97	84.57 \pm 0.57	86.53 \pm 0.79	92.11 \pm 0.35
CAN+Top- k [16]	ResNet-12	67.19 \pm 0.55	80.64 \pm 0.35	73.21 \pm 0.58	84.93 \pm 0.38	-	-	-	-
DPGN [65]	ResNet-12	67.77 \pm 0.32	84.60 \pm 0.43	72.45 \pm 0.51	87.24 \pm 0.39	77.90 \pm 0.50	90.20 \pm 0.40	75.71 \pm 0.47	91.48 \pm 0.33
MCT (instance) [29]	ResNet-12B	78.55 \pm 0.86	86.03 \pm 0.42	82.32 \pm 0.81	87.36 \pm 0.50	85.61 \pm 0.69	90.03 \pm 0.46	-	-
EP [51]	WRN-28-10	70.74 \pm 0.85	84.34 \pm 0.53	78.50 \pm 0.91	88.36 \pm 0.57	-	-	-	-
SIB [17]†	WRN-28-10	70.00 \pm 0.60	79.20 \pm 0.40	72.90	82.80	80.00 \pm 0.60	85.3 \pm 0.40	-	-
SIB+E ³ BM [37]	WRN-28-10	71.40 \pm 0.50	81.20 \pm 0.40	75.60 \pm 0.6	84.30 \pm 0.4	-	-	-	-
LaplacianShot [69]	WRN-28-10	74.86 \pm 0.19	84.13 \pm 0.14	80.18 \pm 0.21	87.56 \pm 0.15	-	-	-	-
PT+MAP [19]*	WRN-28-10	82.88 \pm 0.73	88.78 \pm 0.40	88.15 \pm 0.71	92.32 \pm 0.40	86.91 \pm 0.72	90.50 \pm 0.49	91.37 \pm 0.61	93.93 \pm 0.32
iLPC (ours)	WRN-28-10	83.05 \pm 0.79	88.82 \pm 0.42	88.50 \pm 0.75	92.46 \pm 0.42	86.51 \pm 0.75	90.60 \pm 0.48	91.03 \pm 0.63	94.11 \pm 0.30

Table 9. *Transductive inference state of the art*. *: our reproduction with official code on our datasets. †: *tieredImageNet* as reported by [37].

METHOD	NETWORK	SPLIT	miniIMAGENET	
			1-shot	5-shot
LST [32]	ResNet-12	30/50	70.10 \pm 1.90	78.70 \pm 0.80
LR+ICI [63]	ResNet-12A	30/50	69.66	80.11
MCT (instance) [29]	ResNet-12B	30/50	73.80 \pm 0.70	84.40 \pm 0.50
k -means [50]†	WRN-28-10	100/100	52.35 \pm 0.89	67.67 \pm 0.65
TransMatch [66]	WRN-28-10	100/100	63.02 \pm 1.07	81.06 \pm 0.59
PTN [20]	WRN-28-10	100/100	81.57 \pm 0.94	87.17 \pm 0.58
iLPC (ours)	WRN-28-10	100/100	87.62 \pm 0.67	90.51 \pm 0.36

Table 10. *Semi-supervised few-shot learning state of the art*. †: as reported by [66].

4.7. Comparison with the state of the art

Table 9 and Table 10 compare our iLPC with a larger collection of recent methods on the transductive and semi-supervised settings, respectively. Even when the network and data split appears to be the same, we acknowledge that our results are not directly comparable with any method other than our reproductions. As discussed in the supplementary material, this is due to the very diverse choices made in the bibliography, e.g. versions of network, training settings, versions of datasets, or pre-processing. For instance, ResNet-12 is different than either ResNet-12A or ResNet-12B.

For this reason, we focus on the best result by each method, including ours. Necessarily, methods experimenting with WRN-28-10 have an advantage. Still, at least among those, iLPC performs best by a large margin in both settings, with the closest second best being PT+MAP [19].

5. Conclusion

Our solution is conceptually simple and combines in a unique way ideas that have been successful in problems related to our task at hand. Label propagation exploits the manifold structure of the data, which becomes important in the presence of more data, while still being competitive otherwise. Class balancing provides a strong hint in correcting predictions when certain classes dominate. Label cleaning, originally introduced for learning with noisy labels, is also very successful in cleaning predicted pseudo-labels. Iterative reuse of few pseudo-labels as true labels bypasses the difficulty of single-shot detection of clean examples.

Importantly, reasonable baselines, like predicting pseudo-labels by a classifier or iteratively re-using pseudo-labels without cleaning, fail completely. When compared under fair settings, our iLPC outperforms or is on par with state-of-the-art methods. It is also significantly more robust against feature pre-processing on which other methods rely.

References

- [1] Paul Albert, Diego Ortego, Eric Arazo, Noel E O'Connor, and Kevin McGuinness. Relab: Reliable label bootstrapping for semi-supervised learning. *arXiv preprint arXiv:2007.11866*, 2020.
- [2] A. Antoniou, A. Storkey, and H. Edwards. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*, abs/1711.04340, 2017.
- [3] Eric Arazo, Diego Ortego, Paul Albert, Noel E. O'Connor, and Kevin McGuinness. Unsupervised label noise modeling and loss correction. In *ICML*, 2019.
- [4] Gedas Bertasius, Lorenzo Torresani, Stella X. Yu, and Jianbo Shi. Convolutional random walk networks for semantic image segmentation. In *CVPR*, July 2017.
- [5] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019.
- [6] Luca Bertinetto, Joao F Henriques, Philip HS Torr, and Andrea Vedaldi. Meta-learning with differentiable closed-form solvers. *arXiv preprint arXiv:1805.08136*, 2018.
- [7] Luca Bertinetto, João F Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In *NeurIPS*, 2016.
- [8] Olivier Chapelle, Bernhard Schlkopf, and Alexander Zien. *Semi-Supervised Learning*. The MIT Press.
- [9] Da Chen, Yuefeng Chen, Yuhong Li, Feng Mao, Yuan He, and Hui Xue. Self-supervised learning for few-shot image classification. *arXiv preprint arXiv:1911.06045*, 2019.
- [10] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019.
- [11] Guneet S Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. *arXiv preprint arXiv:1909.02729*, 2019.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [13] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018.
- [14] Bharath Hariharan and Ross Girshick. Low-shot visual recognition by shrinking and hallucinating features. In *CVPR*, 2017.
- [15] 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.
- [16] Ruibing Hou, Hong Chang, MA Bingpeng, Shiguang Shan, and Xilin Chen. Cross attention network for few-shot classification. In *NeurIPS*, 2019.
- [17] Shell Xu Hu, Pablo Garcia Moreno, Yang Xiao, Xi Shen, Guillaume Obozinski, Neil Lawrence, and Andreas Damianou. Empirical bayes transductive meta-learning with synthetic gradients. In *ICLR*, 2019.
- [18] Yuqing Hu, Vincent Gripon, and S. Pateux. Exploiting unsupervised inputs for accurate few-shot classification. *arXiv preprint arXiv:2001.09849*, 2020.
- [19] Yuqing Hu, Vincent Gripon, and Stéphane Pateux. Leveraging the feature distribution in transfer-based few-shot learning. *arXiv preprint arXiv:2006.03806*, 2020.
- [20] H. Huang, Junjie Zhang, Jian Zhang, Qiang Wu, and Chang Xu. Ptn: A poisson transfer network for semi-supervised few-shot learning. In *AAAI*, 2021.
- [21] Jinchu Huang, Lie Qu, Rongfei Jia, and Binqiang Zhao. O2U-Net: A simple noisy label detection approach for deep neural networks. *ICCV*, 2019.
- [22] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *CVPR*, 2019.
- [23] Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. Proxy anchor loss for deep metric learning. In *CVPR*, 2020.
- [24] Philip A Knight. The Sinkhorn-Knopp algorithm: convergence and applications. *SIAM Journal on Matrix Analysis and Applications*, 2008.
- [25] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICML workshop*, 2015.
- [26] Deguang Kong and Chris Ding. Maximum consistency preferential random walks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2012.
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Image classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [29] Seong Min Kye, Hae Beom Lee, Hoirin Kim, and Sung Ju Hwang. Meta-learned confidence for few-shot learning. *arXiv preprint arXiv:2002.12017*, 2020.
- [30] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. 2013.
- [31] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, 2019.
- [32] Xinzhe Li, Qianru Sun, Yaoyao Liu, Qin Zhou, Shibao Zheng, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. In *NeurIPS*, 2019.
- [33] Moshe Lichtenstein, Prasanna Sattigeri, Rogerio Feris, Raja Giryes, and Leonid Karlinsky. TAFSSL: Task-adaptive feature sub-space learning for few-shot classification, 2020.
- [34] Bin Liu, Yue Cao, Yutong Lin, Qi Li, Zheng Zhang, Ming-sheng Long, and Han Hu. Negative margin matters: Understanding margin in few-shot classification. In *ECCV*, 2020.
- [35] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *CVPR*, 2019.
- [36] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. *arXiv preprint arXiv:1805.10002*, 2018.
- [37] Yaoyao Liu, Bernt Schiele, and Qianru Sun. An ensemble of epoch-wise empirical bayes for few-shot learning. In *ECCV*, 2020.

- [38] Puneet Mangla, Nupur Kumari, Abhishek Sinha, Mayank Singh, Balaji Krishnamurthy, and Vineeth N Balasubramanian. Charting the right manifold: Manifold mixup for few-shot learning. In *WACV*, 2020.
- [39] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. *arXiv preprint arXiv:1707.03141*, 2017.
- [40] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *ICML*, 2017.
- [41] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- [42] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018.
- [43] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [44] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *JMLR*, 2011.
- [45] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *CVPR*, 2018.
- [46] L. Qiao, Y. Shi, Jia Li, Yaowei Wang, Tiejun Huang, and Yonghong Tian. Transductive episodic-wise adaptive metric for few-shot learning. *ICCV*, 2019.
- [47] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, 2018.
- [48] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *NeurIPS*, 2019.
- [49] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. 2016.
- [50] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. *arXiv preprint arXiv:1803.00676*, 2018.
- [51] Pau Rodríguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. *ECCV*, 2020.
- [52] Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*, 2018.
- [53] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, 2016.
- [54] H Scudder. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 1965.
- [55] Christian Simon, Piotr Koniusz, Richard Nock, and Mehrta Harandi. Adaptive subspaces for few-shot learning. In *CVPR*, 2020.
- [56] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [57] Jiaming Song, Lunjia Hu, Yann Dauphin, M. Auli, and Tengyu Ma. Robust and on-the-fly dataset denoising for image classification. *arXiv preprint arXiv:2003.10647*, 2020.
- [58] 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.
- [59] 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.
- [60] Eleni Triantafillou, Richard Zemel, and Raquel Urtasun. Few-shot learning through an information retrieval lens. In *NeurIPS*, 2017.
- [61] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, 2016.
- [62] Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *CVPR*, 2019.
- [63] Yikai Wang, C. Xu, Chen Liu, Liyong Zhang, and Yanwei Fu. Instance credibility inference for few-shot learning. *CVPR*, 2020.
- [64] Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. In *CVPR*, 2018.
- [65] Ling Yang, Liangliang Li, Zilun Zhang, Xinyu Zhou, Erjin Zhou, and Yu Liu. Dpgn: Distribution propagation graph network for few-shot learning. In *CVPR*, 2020.
- [66] Zhongjie Yu, L. Chen, Zhongwei Cheng, and Jiebo Luo. Transmatch: A transfer-learning scheme for semi-supervised few-shot learning. *CVPR*, 2020.
- [67] Dengyong Zhou, Olivier Bousquet, Thomas Navin Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *NIPS*, 2003.
- [68] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical report, 2002.
- [69] Imtiaz Ziko, Jose Dolz, Eric Granger, and Ismail Ben Ayed. Laplacian regularized few-shot learning. In *International Conference on Machine Learning*, pages 11660–11670. PMLR, 2020.
- [70] Luisa Zintgraf, Kyriacos Shiarli, Vitaly Kurin, Katja Hofmann, and Shimon Whiteson. Fast context adaptation via meta-learning.