

# A Strong Baseline for Semi-Supervised Incremental Few-Shot Learning

Linglan Zhao<sup>\*1</sup>

llzhao@sjtu.edu.cn

Dashan Guo<sup>2</sup>, Yunlu Xu<sup>†2</sup>, Liang Qiao<sup>2</sup>

{guodashan,xuyunlu,qiaoliang6}@hikvision.com<sup>2</sup>

Zhanzhan Cheng<sup>2</sup>, Shiliang Pu<sup>2</sup>, Yi Niu<sup>2</sup>

{chengzhanzhan,pushiliang,niuyl}@hikvision.com

Xiangzhong Fang<sup>1</sup>

xzfang@sjtu.edu.cn

<sup>1</sup> Dept. of Electronic Engineering

Shanghai Jiao Tong University

Shanghai, China

<sup>2</sup> Hikvision Research Institute

Hangzhou, China

## Abstract

Few-shot learning (FSL) aims to learn models that generalize to novel classes with limited training samples. Recent works advance FSL towards a scenario where unlabeled examples are also available and propose semi-supervised FSL methods. Another line of methods also cares about the performance of base classes in addition to the novel ones and thus establishes the incremental FSL scenario. In this paper, we generalize the above two under a more realistic yet complex setting, named by Semi-Supervised Incremental Few-Shot Learning (S<sup>2</sup>I-FSL). To tackle the task, we propose a novel paradigm containing two parts: (1) a well-designed meta-training algorithm for mitigating ambiguity between base and novel classes caused by unreliable pseudo labels and (2) a model adaptation mechanism to learn discriminative features for novel classes while preserving base knowledge using few labeled and all the unlabeled data. Extensive experiments on standard FSL, semi-supervised FSL, incremental FSL, and the firstly built S<sup>2</sup>I-FSL benchmarks demonstrate the effectiveness of our proposed method.

## 1 Introduction

Despite the great success of deep learning models in recent years, modern deep learning approaches rely heavily on abundant labeled data which can be prohibitively expensive to acquire. In contrast, humans can quickly learn novel concepts with only a few samples. To tackle this problem, *Few-Shot Learning (FSL)* [8, 31, 33, 38] methods are proposed to bridge the gap between supervised learning and the intriguing property of human cognition.

The *standard FSL* setting is shown in Fig. 1 (a) that given new tasks of limited supervision, the system turns the prior knowledge from base tasks. Recent works [14, 18, 26, 44] advance the few-shot paradigm towards a scenario where unlabeled examples are available, namely the *semi-supervised FSL* shown in Fig. 1 (b). Exploring extra unlabeled data along with the limited supervised ones makes sense since unlabeled data is much easier to obtain than the labeled. However, the standard or semi-supervised FSL focuses solely on recognizing few novel classes, while ignoring the ability to handle the previously learned base classes.

\* Linglan Zhao did this work when he was an intern at Hikvision Research Institute. †: corresponding author.

© 2021. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

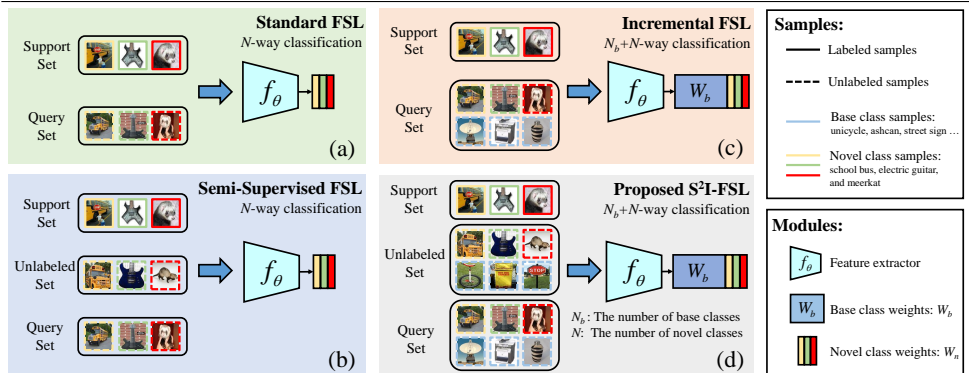


Figure 1: Comparisons between our proposed S<sup>2</sup>I-FSL and other FSL benchmarks.  $N_b$  and  $N$  are the numbers of base and novel classes, respectively. Best viewed in color.

To prevent forgetting prior knowledge for base tasks, *incremental FSL* [9, 23, 27, 43] is built for simultaneously regarding performance on base and novel classes, which shows some realistic significance, as shown in Fig. 1 (c). Nevertheless, there is still room for improvement as only supervised data is used which is hard to adapt to new classes.

From a more realistic perspective, besides few annotated samples, we expect to additionally utilize more easily available unlabeled data containing both base and novel class samples to overcome the performance limitation, especially the performance of novel classes. Therefore, we propose a new benchmark incorporating the above semi-supervised and incremental settings together called *Semi-Supervised Incremental Few-Shot Learning (S<sup>2</sup>I-FSL)*, which generalizes previous assumptions under a more realistic yet complex setting (Fig. 1 (d)).

Intuitively, a straightforward solution for the proposed S<sup>2</sup>I-FSL is to simply combine the existing semi-supervised and incremental FSL approaches, but it does not work well. For detail, as illustrated in Section 4.1, we utilize a simple yet efficient meta-learning-based method as our incremental FSL baseline (denoted as Meta-Inc-Baseline). When extended with unlabeled data in a semi-supervised manner by the mainstreaming methods including weights refinement [18, 19, 26], label propagation [20, 28] and FixMatch [62], unfortunately, it can not perform well. We find that novel class weights refined with incorrectly predicted pseudo-labels will overlap with base class weights and lead to ambiguity among classes in the embedding space. For example, the weight of novel class “school bus” in Fig. 1 (d) is likely to be incorrectly refined with unlabeled samples from base class “unicycle” (both have the *wheel* pattern), “ashcan” (similar color and texture), and “street sign” (two objects often appear in the same image), especially when only a few labeled “school bus” samples are given as in few-shot scenarios. Overall, S<sup>2</sup>I-FSL brings additional challenges in the FSL field: (1) using unlabeled data from both base and novel classes leads easily to ambiguity among classes in the embedding space, showing performance degradation; (2) unbalanced sample number of different classes in unlabeled set instead of a balanced number in each novel class makes it harder to learn the classifier. In this paper, we mainly handle the former challenge and show the comparison effect on the latter one.

Upon the above, we propose a novel S<sup>2</sup>I-FSL paradigm of two phases including two key components, a novel meta-training scheme and a well-designed model adaptation mechanism. In the meta-training phase, the training set is sampled using a mechanism to mimic the semi-supervised scenario during testing. As a result, the model learns to make use of unlabeled data for incremental learning and mitigates confusion between base and novel

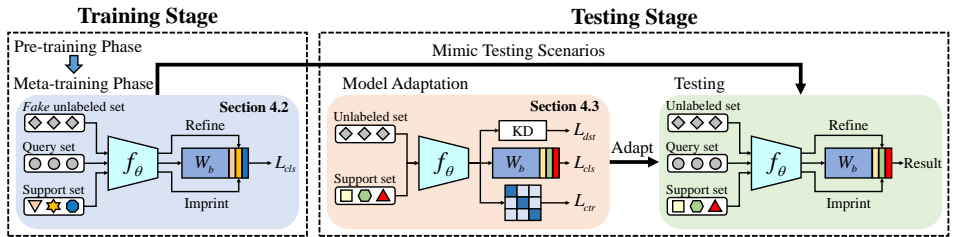


Figure 2: An overview of our proposed method for  $S^2I$ -FSL. The left and right parts illustrate the model training and testing stages, respectively. Best viewed in color.

classes. Considering that the feature extractor is trained on the training set, it may not contain discriminative information for recognizing novel classes unseen during training, a model adaptation scheme is proposed. Concretely, besides standard classification loss on few labeled data, a contrastive learning loss function is applied to unlabeled data to explore the novel class distribution implicitly. In addition, we also use distillation technique to preserve the classification ability on base classes.

To summarize, the contributions of this work are three-fold: (1) A Semi-Supervised Incremental Few-Shot Learning ( $S^2I$ -FSL) benchmark is proposed to generalize previous semi-supervised FSL and incremental FSL under a more realistic and challenging setting; (2) Technically, we propose an efficient meta-training paradigm and a model adaptation mechanism for the newly built task, which can be regarded as a baseline for future researches; (3) Extensive experiments on standard FSL, incremental FSL, semi-supervised FSL, and the  $S^2I$ -FSL benchmark demonstrate the effectiveness of our proposed method.

## 2 Related works

### 2.1 Few-Shot Learning

Few-shot learning (FSL) aims to learn novel categories with limited labeled examples. Previous episodic-training [68] based FSL works can be roughly divided into 4 categories. *Metric learning based* approaches [60, 63, 68] aim at learning the similarity metric between query (*i.e.*, test) samples and support (*i.e.*, training) samples. The goal of *initialization based* FSL methods [0, 8, 29] is to learn good model initialization for fast adaptation to current tasks. *Weight generation based* methods [9, 23, 25] directly predict classification weights for novel classes. Moreover, *hallucination based* approaches [16, 40, 47] learn a generator for data augmentation. However, recent works [6, 6, 10, 65] reveal that a well-trained feature extractor, rather than other complex meta-trained modules, is the key component to achieve state-of-the-art performance. This motivates us to adopt a modified version of [6] as a baseline method (Meta-Inc-Baseline) for incremental FSL and  $S^2I$ -FSL benchmark.

### 2.2 Few-Shot Incremental Learning

Incremental FSL involves training a model to classify novel classes while preserving the capability of handling previously learned classes. There are two main different settings in the literature. The first setting is also called generalized (or dynamic) FSL, which contains two sequences of tasks. The first task contains abundant training samples from base classes, while the second task contains only a few novel class training samples. Methods proposed in this setting include novel class weight generation [9, 15, 23, 30], attention-based regularization [27] and task adaptive features extraction [22, 43]. The second setting [3, 21, 54, 46] shares the same core idea with the first one but considers several sequences of tasks instead

of two. We build S<sup>2</sup>I-FSL mainly based on the first setting since extra unlabeled data are not available in the second setting. In addition, the second setting uses unique data splits which makes it harder to compare to other standard/transductive/semi-supervised FSL works.

### 2.3 Semi-Supervised Few-Shot Learning

Semi-supervised/transductive FSL works take additional unlabeled data into account to boost the performance of FSL. *Pseudo-labeling based* methods and *graph-based* methods are two main lines of efforts. Pseudo-labeling based methods consist of prototype refinement [14, 18, 19, 26], self-training [17, 39] and entropy minimization [2, 2]. As for graph-based approaches [24, 21, 28, 41], graph models are constructed to propagate information from labeled data to unlabeled data. Among these approaches, we mainly focus on the utilization of prototype refinement for its effectiveness and extensiveness.

## 3 Task Formulation

In this section, we introduce the detailed settings for the proposed S<sup>2</sup>I-FSL benchmark.

**Dataset Splits.** For evaluating classification performance on both base and novel classes, we follow the dataset splits in incremental FSL [9, 27, 43]. Concretely,  $\mathcal{D}_{\text{base}}$  and  $\mathcal{D}_{\text{novel}}$  are two datasets for base classes  $\mathcal{C}_{\text{base}}$  and novel classes  $\mathcal{C}_{\text{novel}}$  respectively, and  $\mathcal{C}_{\text{base}} \cap \mathcal{C}_{\text{novel}} = \emptyset$ .  $\mathcal{D}_{\text{base}}$  is comprised of  $\mathcal{D}_{\text{base/train}}$ ,  $\mathcal{D}_{\text{base/val}}$  and  $\mathcal{D}_{\text{base/test}}$  as base class training, validation and testing sets ( $|\mathcal{C}_{\text{base}}| = N_b$ ).  $\mathcal{D}_{\text{novel}}$  consists of three splits of disjoint novel classes:  $\mathcal{D}_{\text{novel/train}}$ ,  $\mathcal{D}_{\text{novel/val}}$  and  $\mathcal{D}_{\text{novel/test}}$  for meta-training, validation and testing, respectively.

**Training Stage.** As shown in the left part of Fig. 2, the training stage is split into *pre-training phase* and *meta-training phase*. In the pre-training phase, feature extractor  $f_\theta$  and the base class weights  $W_b$  are trained on  $\mathcal{D}_{\text{base/train}}$  as in standard supervised learning. After that, in the meta-training phase, the model is further trained using both base class training set  $\mathcal{D}_{\text{base/train}}$  and novel class training set  $\mathcal{D}_{\text{novel/train}}$  in an episodic training manner [58].

**Testing Stage.** Following the evaluation protocol in [9, 27, 43], average accuracy on sampled incremental FSL episodes (*i.e.*, tasks) is reported. For each test episode,  $N$  novel categories are randomly chosen from  $\mathcal{D}_{\text{novel/test}}$  ( $\mathcal{D}_{\text{novel/val}}$  for validation), which are relabeled with new labels from  $N_b + 1$  to  $N_b + N$ . After that,  $K$  per-class training samples are sampled from these  $N$  classes ( $N$ -way  $K$ -shot) to construct a support set  $\mathcal{S}$ . The query set  $\mathcal{Q}$  consists of samples from both novel and base classes, *i.e.*, additional examples from the selected  $N$  novel classes and base class test examples from  $\mathcal{D}_{\text{base/test}}$  ( $\mathcal{D}_{\text{base/val}}$  for validation). The main difference between S<sup>2</sup>I-FSL and incremental FSL is that unlabeled data is available in S<sup>2</sup>I-FSL. Two different settings are considered: a *transductive* setting where the unlabeled set  $\mathcal{U}$  is exactly the query set  $\mathcal{Q}$ ; a *semi-supervised* setting where  $\mathcal{U}$  contains extra unlabeled samples from both base classes and novel classes.

## 4 Methodology

### 4.1 Preliminary

Since Meta-Baseline [9] provides an effective yet strong baseline for standard FSL, we extend it for incremental setting, noted as Meta-Inc-Baseline. During the *pre-training phase*, the feature extractor  $f_\theta(\mathbf{x}) \in \mathbb{R}^d$  is trained using a cosine classifier on  $\mathcal{D}_{\text{base/train}}$  as in standard supervised training. The conditional probability of sample  $\mathbf{x}$  belonging to class  $k$  is:

$$P(y = k \mid \mathbf{x}; W_b) = \frac{\exp(\gamma \cos(f_\theta(\mathbf{x}), w_k))}{\sum_j \exp(\gamma \cos(f_\theta(\mathbf{x}), w_j))}, \quad (1)$$

where  $W_b = [w_1, \dots, w_{N_b}] \in \mathbb{R}^{d \times N_b}$  is the base class weights and  $\gamma$  is a learnable scalar.

In the *testing stage*, novel class weights can be constructed as class means of the support set  $\mathcal{S}$ , i.e., the novel class weights  $W_n = [p_1, \dots, p_N] \in \mathbb{R}^{d \times N}$  are comprised of prototypes [43] of each category where  $\mathcal{S}_j = \{(\mathbf{x}, y) \mid (\mathbf{x}, y) \in \mathcal{S}, y = N_b + j\}$ :

$$p_j = \frac{1}{|\mathcal{S}_j|} \sum_{(\mathbf{x}, y) \in \mathcal{S}_j} f_\theta(\mathbf{x}). \quad (2)$$

For extending Meta-Baseline to handle both base and novel classes, the overall classification weights  $W$  is constructed by concatenating the base class weights  $W_b$  and novel class weights  $W_n$ , i.e.  $W = [W_b, W_n] \in \mathbb{R}^{d \times (N_b + N)}$ . Finally, the model makes predictions on query set  $\mathcal{Q}$  containing both base and novel class samples based on Eq. 1 with  $W_b$  replaced with  $W$ .

After pre-training is a *meta-training phase*. Since the model is evaluated on incremental FSL tasks, instead of only optimizing  $f_\theta$  and discarding  $W_b$  [6], we meta-train both  $f_\theta$  and  $W_b$  by sampling incremental FSL tasks from  $\mathcal{D}_{\text{base/train}}$  and  $\mathcal{D}_{\text{novel/train}}$  as in Algorithm 1.

---

### Algorithm 1 Meta-training algorithm for Meta-Inc-Baseline.

---

**Input:** pre-trained  $f_\theta$  and  $W_b$ , learning rate for  $\theta$  and  $W_b$  respectively:  $\eta_1$  and  $\eta_2$ .

**Output:** Meta-trained feature extractor  $f_\theta$  and base class weights  $W_b$ .

- 1: **while** not done **do**
  - 2:    $\{\mathcal{S}, \mathcal{Q}_n\} \leftarrow$  Sample support and query set from  $\mathcal{D}_{\text{novel/train}}$
  - 3:    $\{\mathcal{Q}_b\} \leftarrow$  Sample query set from  $\mathcal{D}_{\text{base/train}}$
  - 4:    $\mathcal{Q} = \mathcal{Q}_b \cup \mathcal{Q}_n$  // **Construct an incremental FSL episode**
  - 5:   Construct overall classification weight  $W = [W_b, W_n]$  using Eq. 2
  - 6:   Get predictions  $\hat{Y}_q \in \mathbb{R}_+^{n_q \times (N_b + N)}$  on  $\mathcal{Q}$  ( $n_q = |\mathcal{Q}|$ ) using Eq. 1 with  $W_b \leftarrow W$
  - 7:   Calculate cross-entropy loss  $\mathcal{L}_{cls}(\hat{Y}_q, Y_q)$  based on  $\hat{Y}_q$  and ground truth labels  $Y_q$
  - 8:   Update:  $\theta \leftarrow \theta - \eta_1 \nabla_\theta \mathcal{L}_{cls}$ ,  $W_b \leftarrow W_b - \eta_2 \nabla_{W_b} \mathcal{L}_{cls}$
  - 9: **end while**
- 

## 4.2 Prototype Refinement with Fake Unlabeled Data

Upon Meta-Inc-Baseline, we further generalize to  $\text{S}^2\text{I-FSL}$  in Section 4.2 and Section 4.3. Since prototype refinement using pseudo labels has been proved effective in semi-supervised methods [13, 18, 19, 26], an intuitive solution to extend Meta-Inc-Baseline to  $\text{S}^2\text{I-FSL}$  is to refine the novel class weights  $W_n$  with unlabeled data. Meanwhile, base class weights  $W_b$  trained on abundant base class samples remain unchanged.

Concretely, predictions  $\hat{Y}_u \in \mathbb{R}_+^{n_u \times (N_b + N)}$  on  $\mathcal{U}$  ( $n_u = |\mathcal{U}|$ ) can be obtained as in line 5-6 of Algorithm 1 similar to  $\hat{Y}_q$ . Since  $\mathcal{U}$  contains unlabeled samples from both base and novel classes, samples predicted as base classes should be filtered out as they are viewed as distractors in the refinement process:  $\hat{Y}' = \hat{Y}_u[:, N_b : N_b + N] \in \mathbb{R}_+^{n_u \times N}$  where  $[:, :]$  denotes slicing operation. Based on  $\hat{Y}'$ , novel class prototypes  $\{p'_j\}_{j=1}^N$  can be re-estimated :

$$p'_j = \frac{\sum_{i=1}^{n_u} \hat{Y}'_{ij} f_\theta(\mathbf{x}_i) + \sum_{(\mathbf{x}_k, y_k) \in \mathcal{S}_j} f_\theta(\mathbf{x}_k)}{\sum_{i=1}^{n_u} \hat{Y}'_{ij} + |\mathcal{S}_j|}. \quad (3)$$

Then we progressively refine  $W_n = [p_1, \dots, p_N] \in \mathbb{R}^{d \times N}$ :

$$p_j \leftarrow \alpha \cdot p'_j + (1 - \alpha) \cdot p_j, \quad (4)$$

where  $\alpha \in (0, 1]$ . After iterating the above process of Eq. 3-4 for  $n_{\text{steps}}$  loops, query samples are evaluated with  $W = [W_b, W_n]$  where  $W_n$  is the refined novel class weights.

However, as discussed in Section 1 and verified in Section 5, this straightforward modification leads to little improvement or even obvious degradation. It is because that novel class weights refined with incorrectly predicted pseudo-labels will be confused with similar base

classes. For utilization of unlabeled samples and mitigating ambiguity, a novel meta-training algorithm with *fake*<sup>1</sup> unlabeled examples is proposed in Algorithm 2 as a substitution for Algorithm 1. Concretely, compared to labeled support set  $\mathcal{S}$  or query set  $\mathcal{Q}$  whose labels are used in the loss function, the label information of sampled unlabeled set  $\mathcal{U}$  in each training episode is not used to mimic the testing scenario. By default,  $n_{\text{steps}} = 1$ ,  $\alpha = 1.0$  in training.

---

### Algorithm 2 Meta-training with *fake* unlabeled data.

---

**Input:** pre-trained  $f_\theta$  and  $W_b$ , learning rate for  $\theta$  and  $W_b$ :  $\eta_1$  and  $\eta_2$ . refinement parameters:  $n_{\text{steps}} = 1$ ,  $\alpha = 1.0$ .

**Output:** Meta-trained feature extractor  $f_\theta$  and base class weights  $W_b$ .

```

1: while not done do
2:    $\{\mathcal{S}, \mathcal{Q}_n, \mathcal{U}_n\} \leftarrow$  Sample support, query and unlabeled set from  $\mathcal{D}_{\text{novel/train}}$ 
3:    $\{\mathcal{Q}_b, \mathcal{U}_b\} \leftarrow$  Sample query and unlabeled set from  $\mathcal{D}_{\text{base/train}}$ 
4:    $\mathcal{Q} = \mathcal{Q}_b \cup \mathcal{Q}_n, \mathcal{U} = \mathcal{U}_b \cup \mathcal{U}_n$  // Construct an incremental FSL episode with fake unlabeled set
5:   Get refined  $W_n$  using Eq. 3-4 based on  $\mathcal{U}$ , and construct classification weight  $W = [W_b, W_n]$ 
6:   Get predictions  $\hat{Y}_q \in \mathbb{R}_+^{n_q \times (N_b + N)}$  on  $\mathcal{Q}$  ( $n_q = |\mathcal{Q}|$ ) using Eq. 1 with  $W_b \leftarrow W$ 
7:   Calculate cross-entropy loss  $\mathcal{L}_{cls}(\hat{Y}_q, Y_q)$  based on  $\hat{Y}_q$  and ground truth labels  $Y_q$ 
8:   Update:  $\theta \leftarrow \theta - \eta_1 \nabla_\theta \mathcal{L}_{cls}, W_b \leftarrow W_b - \eta_2 \nabla_{W_b} \mathcal{L}_{cls}$ 
9: end while

```

---

## 4.3 Model Adaptation without Forgetting

As described in Section 1, feature extractor  $f_\theta$  trained on the training set may not contain discriminative information for novel classes during testing [6, 14, 36]. However, as observed in few-shot incremental literature [9, 27, 43], directly finetuning models on few labeled novel class samples will cause not only overfitting to the novel class but also catastrophic forgetting of the old classes. Fortunately, since additional unlabeled data is available,  $\{f_\theta, W_b\}$  can be further adapted to the current test episode  $\mathcal{T} = \{\mathcal{S}, \mathcal{Q}, \mathcal{U}\}$  in the testing stage.

As shown in the right part of Fig. 2, for learning better representations for novel classes, we explore both labeled support set  $\mathcal{S}$  using cross-entropy loss  $\mathcal{L}_{cls}$  for classification, and the unlabeled set  $\mathcal{U}$  using a contrastive loss [4]  $\mathcal{L}_{ctr}$ :

$$\mathcal{L}_{ctr} = -\frac{1}{2B} \sum_{i,j=1}^B \log \frac{\exp(\cos(f_i, f_j) / \tau_1)}{\sum_{k \neq i} \exp(\cos(f_i, f_k) / \tau_1)}, \quad (5)$$

where  $B$  is the batch size to sample batches from  $\mathcal{U}$ ,  $\tau_1$  is a temperature parameter and  $(f_i, f_j)$  is a positive feature pair (different augmentations from the same image) based on  $f_\theta$ .

For preserving the knowledge learned from base classes, we adopt a distillation [42] loss  $\mathcal{L}_{dst}$  on base class predictions using unlabeled data  $\mathcal{U}$ :

$$\mathcal{L}_{dst} = -\frac{1}{B} \sum_{i=1}^B \sum_{k=1}^{N_b} \bar{z}_{ik} \log z_{ik}, \quad (6)$$

where  $\bar{z}_{ik}$  and  $z_{ik}$  are  $\tau_2$  softened predictions of the  $i$ -th sample belonging to class  $k$  from previous meta-trained parameters  $\{\bar{\theta}, \bar{W}_b\}$  and current adapted ones  $\{\theta, W_b\}$ , respectively:

$$\bar{z}_{ik} = \frac{\exp(\gamma \cos(f_{\bar{\theta}}(\mathbf{x}_i), \bar{w}_k) / \tau_2)}{\sum_{j=1}^{N_b} \exp(\gamma \cos(f_{\bar{\theta}}(\mathbf{x}_i), \bar{w}_j) / \tau_2)}, \quad z_{ik} = \frac{\exp(\gamma \cos(f_\theta(\mathbf{x}_i), w_k) / \tau_2)}{\sum_{j=1}^{N_b} \exp(\gamma \cos(f_\theta(\mathbf{x}_i), w_j) / \tau_2)}. \quad (7)$$

The overall loss function is formulated as:  $\mathcal{L} = w_{cls} \cdot \mathcal{L}_{cls} + w_{ctr} \cdot \mathcal{L}_{ctr} + w_{dst} \cdot \mathcal{L}_{dst}$  where  $w_{cls}$ ,  $w_{ctr}$  and  $w_{dst}$  (hyperparameters in Appendix C) control the tradeoff between each term.

---

<sup>1</sup>We use the terminology *fake* unlabeled since a subset is randomly sampled from the supervised training set in each meta-training episode, but its labels are not used to simulate the testing scenario.



## 5 Experiments

### 5.1 Experiment Setup

**Datasets.** We follow the common incremental FSL setting [9, 27, 43], and conduct experiments on *mini*-ImageNet [48] and *tiered*-ImageNet [26] with 64 and 200<sup>2</sup> base classes, respectively. More details of dataset statistics can be found in Appendix A.

**Evaluation Protocol.** The performance is evaluated using the average accuracy and corresponding 95% confidence interval over 600 randomly sampled episodes with a *fixed seed*, which makes results more stable and reliable. In incremental FSL and S<sup>2</sup>I-FSL, we report joint accuracy  $\text{Acc}_{\text{all/all}}$  (Acc. for short) which is the accuracy of classifying all the test samples to all (both base and novel) classes. To measure the performance gap between joint accuracy and individual accuracy, following [27, 43], metric  $\Delta$  is also reported. Concretely,  $\Delta = (\Delta_b + \Delta_n)/2$  is computed as the mean of performance degradation on base and novel classes.  $\Delta_b$  ( $\Delta_n$ ) is the performance degradation between the accuracy of classifying base (novel) class test samples to all classes  $\text{Acc}_{b/\text{all}}$  ( $\text{Acc}_{n/\text{all}}$ ) and the accuracy of classifying base (novel) class test samples to only base (novel) classes  $\text{Acc}_{b/b}$  ( $\text{Acc}_{n/n}$ ):

$$\Delta_b = \text{Acc}_{b/\text{all}} - \text{Acc}_{b/b}, \Delta_n = \text{Acc}_{n/\text{all}} - \text{Acc}_{n/n}. \quad (8)$$

**Implementation Details.** In most of the experiments, ResNet12 [11, 22] (detailed in Appendix B) is adopted as the backbone  $f_\theta$  to have fair comparisons between other methods. Unless specified in S<sup>2</sup>I-FSL, unlabeled base and novel class samples are maintained in equal proportion. Concretely, in each test episode, we sample 5 novel classes with 1/5 per-class training samples (5-way 1/5-shot) for support set  $\mathcal{S}$ . Query set  $\mathcal{Q}$  contains  $5 \times 15$  novel class test samples and 75 examples sampled uniformly from the base class test set [27, 43]. After adopting the semi-supervised FSL [11, 18, 59] setting, unlabeled set  $\mathcal{U}$  in S<sup>2</sup>I-FSL contains  $5 \times 30 / 5 \times 50$  novel class samples and 150 / 250 base class samples for 1/5-shot setting. We leave more details of hyper-parameter settings in Appendix C.

### 5.2 Experimental Results

**Incremental FSL.** We first conduct experiments to evaluate the robustness of the proposed Meta-Inc-Baseline on the incremental FSL classification task. The upper parts of Table 1 and Table 2 show the results (denoted as *Inductive*) on the *mini*-ImageNet and *tiered*-ImageNet, respectively. Compared to other approaches, our Meta-Inc-Baseline is able to achieve much higher joint accuracy and diminish the gap between joint accuracy and individual accuracy. The direct performance enhancement is mainly attributed to the extension of the simple but strong baseline of Meta-baseline [6]. The detailed verification is described in Appendix D. Moreover, the effectiveness of Meta-Inc-Baseline on incremental FSL benchmark with multi sessions [44] is also verified in Appendix E.

**S<sup>2</sup>I-FSL.** The proposed method is then generalized into two settings where unlabeled data is available: the *Transductive* setting where unlabeled set  $\mathcal{U}$  is exactly the query set  $\mathcal{Q}$ , and the *Semi-supervised* setting where extra unlabeled set  $\mathcal{U}$  is available. Besides our proposed method, extensions with prevailing semi-supervised FSL methods such as label propagation [20] (a graph model is constructed with base and novel class weights as labeled vertices and unlabeled samples as unlabeled vertices), prototype refinement [18] and semi-supervised learning approach FixMatch [52] are also evaluated for comparison.

As shown in the bottom parts of Table 1 and Table 2, label propagation [20] is not suitable for the more challenging task since it always results in lower accuracy and degrades  $\Delta$  metric.

<sup>2</sup>Due to the dataset splits in incremental FSL [27, 43], the number of base classes is 200 instead of 351.

Method	$f_\theta$	Type	64+5-way 1-shot		64+5-way 5-shot		
			Acc.	$\Delta$	Acc.	$\Delta$	
Imprint [12]	Res10	Incremental FSL (Inductive)	41.34±0.54	-23.79	46.34±0.54	-25.25	
LwoF [9]	Res10		49.65±0.64	-14.47	59.66±0.55	-12.35	
Attractor [12]	Res10		54.95±0.30	-11.84	63.04±0.30	-10.66	
TapNet [12]	Res12 <sup>‡</sup>		54.38±0.59	-12.88	64.02±0.51	-10.98	
XtarNet [12]	Res12 <sup>‡</sup>		54.96±0.61	-13.36	64.88±0.47	-10.41	
Meta-Inc-Baseline (ours)	Res12 <sup>‡</sup>		59.27±0.43	-11.30	69.74±0.36	-9.90	
LCwoF <i>unlim</i> [12]	Res12		58.16±N/A	-N/A	66.88±N/A	-N/A	
Meta-Inc-Baseline (ours)	Res12		<b>60.65±0.42</b>	<b>-11.23</b>	<b>71.52±0.35</b>	<b>-9.30</b>	
Meta-Inc-Baseline + Graph [12]	Res12		S <sup>2</sup> I-FSL (Transductive)	55.02±0.41	-18.32	65.71±0.38	-15.57
Meta-Inc-Baseline + FixMatch [12]	Res12			57.14±0.46	-13.24	72.07±0.36	-9.33
Meta-Inc-Baseline + PR [12]	Res12	62.74±0.54		-13.57	72.60±0.37	-9.50	
Our proposed method	Res12	<b>68.43±0.54</b>		<b>-8.31</b>	<b>74.97±0.35</b>	<b>-7.95</b>	
Meta-Inc-Baseline + Graph [12]	Res12	S <sup>2</sup> I-FSL (Semi-supervised)	50.41±0.36	-23.53	62.82±0.38	-18.04	
Meta-Inc-Baseline + FixMatch [12]	Res12		56.60±0.46	-13.59	72.00±0.38	-9.49	
Meta-Inc-Baseline + PR [12]	Res12		64.53±0.54	-12.46	73.48±0.35	-9.00	
Our proposed method	Res12		<b>70.33±0.51</b>	<b>-8.07</b>	<b>75.91±0.33</b>	<b>-7.79</b>	

<sup>‡</sup>: ResNet12 backbone with 256 output channels as [12, 12]. PR: directly Prototype Refinement using unlabeled data.

Table 1: Incremental FSL and S<sup>2</sup>I-FSL results on *mini*-ImageNet.

While FixMatch [12] can get about 0.5% performance gain in 5-shot case, the accuracy decreases significantly in 1-shot case. It is because that FixMatch relies heavily on the quality of pseudo-labels that cannot be obtained in few-shot scenarios especially the 1-shot case. Moreover, the direct combination of strong baseline and prototype refinement [12] can get a limited improvement on *mini*-ImageNet but leads to obvious performance degradation on *tiered*-ImageNet. Since *tiered*-ImageNet contains more base classes, confusion between base and novel classes caused by unreliable pseudo labels is prone to happen.

In contrast, our proposed method can obtain significant performance gains on both joint accuracy and resistance to forgetting consistently. In *Transductive* setting, compared to Meta-Inc-Baseline + PR, we get 5.7%/12.0% 1-shot classification performance gains and 5.3%/10.1%  $\Delta$  metric improvements on the two datasets, respectively. It is because that our proposed method can alleviate confusion between base and novel classes, and is well adapted to the current test episode to learn discriminative features for novel classes. In *Semi-Supervised* setting, our method outperforms Meta-Inc-Baseline by about 10% in 1-shot classification on both benchmarks, which shows the superiority of utilizing the unlabeled set.

For better understanding the challenges of S<sup>2</sup>I-FSL, we report detailed results in Table 3. Although all the methods can improve  $\text{Acc}_{n/n}$  metric which is the only focus of standard/transductive/semi-supervised FSL methods, other competitors obtain decreased joint accuracy or only marginal overall improvement. This indicates that S<sup>2</sup>I-FSL is a more difficult task and poses challenges that have not been solved by existing works. We emphasize that there is a trade-off between  $\text{Acc}_{b/all}$  and  $\text{Acc}_{n/all}$ , and balancing between them is significant for the final result. For example, although  $\text{Acc}_{b/all}$  increases using graph models [12],  $\text{Acc}_{n/all}$  decreases dramatically which results in degraded joint accuracy. As mentioned in Section 1, there is an imbalance between base and novel classes. Since the number of base classes is larger than that of novel classes, label information from base vertices ( $N_b = 200$ ) is more likely propagated to unlabeled vertices which decreases  $\text{Acc}_{n/all}$ . The marginal improvement of FixMatch [12] in 5-shot case is attributed to that the increase in  $\text{Acc}_{n/all}$  slightly outweighs the decrease in  $\text{Acc}_{b/all}$ . However, this limited improvement is unstable, since it gets degraded performance in more challenging 1-shot scenarios on both datasets. As for prototype refinement [12],  $\text{Acc}_{b/all}$  degrades significantly due to the confusion between base and novel classes. Concretely, some base class unlabeled samples are incorrectly predicted



as novel classes and are used to refine novel weights. Base class testing samples may be predicted as novel if the corresponding weight is refined with the similar but incorrect base class unlabeled samples. For example, as discussed in Section 1, novel class weight “school bus” in Fig. 1 (d) is likely to be incorrectly refined with unlabeled samples from base class “unicycle” (same *wheel* pattern), “ashcan” (similar color and texture), and “street sign” (often appear in the same image) which causes confusion. By contrast, our method can boost both  $\text{Acc}_{b/all}$  and  $\text{Acc}_{n/all}$ , which accounts for the obvious improvement on joint accuracy.

Method	$f_\theta$	Type	200+5-way 1-shot		200+5-way 5-shot	
			Acc.	$\Delta$	Acc.	$\Delta$
Imprint [10]	Res18	Incremental FSL (Inductive)	40.83±0.45	-22.29	53.87±0.48	-17.18
LwoF [9]	Res18		53.42±0.56	-9.59	63.22±0.52	-7.27
Attractor [10]	Res18		56.11±0.33	<b>-6.11</b>	65.52±0.31	<b>-4.48</b>
TapNet [10]	Res18		55.42±0.59	-8.09	65.38±0.53	-6.37
XtarNet [10]	Res18		56.57±0.60	-8.02	66.52±0.49	-6.31
Meta-Inc-Baseline (ours)	Res18		<b>63.16±0.51</b>	-7.53	<b>73.28±0.40</b>	-5.35
LCwoF <i>unlim</i> [10]	Res12		58.82±N/A	-N/A	66.61±N/A	-N/A
Meta-Inc-Baseline (ours)	Res12		61.92±0.50	-7.87	72.98±0.38	-5.39
Meta-Inc-Baseline + Graph [10]	Res12	$S^2$ I-FSL (Transductive)	59.73±0.56	-11.25	70.18±0.41	-8.58
Meta-Inc-Baseline + FixMatch [10]	Res12		57.49±0.59	-8.12	73.34±0.38	-5.46
Meta-Inc-Baseline + PR [10]	Res12		57.99±0.61	-15.58	72.08±0.39	-7.47
Our proposed method	Res12		<b>70.03±0.57</b>	<b>-5.41</b>	<b>75.73±0.37</b>	<b>-4.53</b>
Meta-Inc-Baseline + Graph [10]	Res12	$S^2$ I-FSL (Semi-supervised)	51.87±0.48	-19.71	65.86±0.41	-12.63
Meta-Inc-Baseline + FixMatch [10]	Res12		56.56±0.59	-8.37	73.22±0.39	-5.57
Meta-Inc-Baseline + PR [10]	Res12		59.35±0.58	-14.55	72.68±0.39	-6.93
Our proposed method	Res12		<b>71.64±0.56</b>	<b>-4.97</b>	<b>76.24±0.36</b>	<b>-4.21</b>

Table 2: Incremental FSL and  $S^2$ I-FSL results on *tiered-ImageNet*.

Method (Semi-supervised)	200+5-way 5-shot					
	Acc.	$\Delta$	$\text{Acc}_{b/all}$	$\text{Acc}_{n/all}$	$\text{Acc}_{b/b}$	$\text{Acc}_{n/n}$
Meta-Inc-Baseline (baseline)	72.98	-5.39	65.91	80.05	73.36	83.39
Meta-Inc-Baseline + Graph [10]	65.86 ↓↓	-12.63 ↓↓	70.65 ↑↑	61.07 ↓↓	71.28 ↓	85.70 ↑↑
Meta-Inc-Baseline + FixMatch [10]	73.22 ↑	-5.57	64.13 ↓	82.30 ↑↑	73.10	84.47 ↑
Meta-Inc-Baseline + PR [10]	72.68 ↓	-6.93 ↓	60.32 ↓↓	85.04 ↑↑	73.36	85.86 ↑↑
Our proposed method	<b>76.24 ↑↑</b>	<b>-4.21 ↑</b>	67.58 ↑↑	84.90 ↑↑	73.08	<b>87.82 ↑↑</b>

↑(↑↑): Certain metric increases (significantly). ↓(↓↓): Certain metric decreases (significantly).

Table 3: Detailed evaluations of  $S^2$ I-FSL on *tiered-ImageNet* 5-shot setting.

**Semi-Supervised/Transductive FSL.** Our model trained for  $S^2$ I-FSL is also directly applied to regular semi-supervised and transductive FSL benchmarks. Here, we only need to perform classification on novel classes where unlabeled data solely contains novel class samples. As shown in Table 4, our method consistently outperforms other approaches using the same or deeper backbones. It can be attributed to the proposed Algorithm 2 for effectively utilizing unlabeled samples and the model adaptation mechanism for learning discriminative features of novel classes. Experiments on *tiered-ImageNet* are provided in Appendix F.

Method	$f_\theta$	Transductive FSL		Method	$f_\theta$	Semi-supervised FSL	
		1-shot	5-shot			1-shot	5-shot
TPN [10]	Conv4	55.51±0.86	69.86±0.65	Soft k-M [10]	Conv4	50.41±0.31	64.39±0.24
TEAM [10]	Res12	60.07±N/A	75.90±N/A	LST [10]	Res12	70.10±1.90	78.70±0.80
BD-CSPN [10]	Res12	65.94±N/A	79.23±N/A	LR+ICI [10]	Res12	71.11±1.15	81.25±0.69
LR+ICI [10]	Res12	68.70±1.10	79.69±0.66	EPNet [10]	Res12	75.36±1.01	84.07±0.60
DPGN [10]	Res12	67.77±0.32	84.60±0.43	TransMatch [10]	WRN28	63.02±1.07	82.24±0.59
Completion [10]	Res12	73.13±0.85	82.06±0.54	ICA+MSP [10]	DenseNet	78.55±0.25	85.41±0.13
Our proposed	Res12	<b>77.37±0.85</b>	<b>86.02±0.41</b>	Our proposed	Res12	<b>79.12±0.83</b>	<b>87.16±0.38</b>

Table 4: Regular semi-supervised/transductive FSL results on *mini-ImageNet*.

### 5.3 Ablation Studies

**Effectiveness of each component.** We present detailed results of one-by-one employment of different components in Table 5. As discussed in Table 3, directly applying prototype refinement results in decreased joint accuracy and  $\Delta$  metric, which is caused by the confusion between base and novel classes. With the help of our proposed meta-training Algorithm 2, the severe confusion is alleviated and about 11.0%/3.1% absolute performance gains are obtained in 1/5-shot cases. Moreover, model adaptation mechanism that learns discriminative features for novel classes yields another 1.3%/0.4% improvement.

Method (Semi-supervised)	Proto refine	Fake unlabel	Model adapt	200+5-way 1-shot			200+5-way 5-shot		
				Acc.	$\Delta$	Acc <sub>n/n</sub>	Acc.	$\Delta$	Acc <sub>n/n</sub>
Meta-Inc-Baseline				61.92	-7.87	67.79	72.98	-5.39	83.39
Meta-Inc-Baseline + PR	✓			59.35	-14.55	75.41	72.68	-6.93	85.86
Our method w/o adaptation	✓	✓		70.34	-5.90	79.82	75.82	-4.49	87.21
Our full method	✓	✓	✓	<b>71.64</b>	<b>-4.97</b>	<b>80.91</b>	<b>76.24</b>	<b>-4.21</b>	<b>87.82</b>

Table 5: Ablation study of each component proposed for S<sup>2</sup>I-FSL on *tiered-ImageNet*.

**Visualization results.** The above observations are also confirmed by t-SNE [37] visualizations. Fig. 3 shows a 200+5-way 1-shot semi-supervised incremental test episode on *tiered-ImageNet* without (left) and with (right) our proposed components. For clarity, 5 base classes and 5 novel classes are randomly chosen and features of 30 per-class test samples are considered. The blue novel class prototype (triangle) which was confused with green and orange base classes is more separable from them using Algorithm 2, and novel class features (blue and cyan, magenta and black) become more distinguishable after model adaptation.

**Studies on unlabeled set  $\mathcal{U}$ .** To observe the performance on the second challenge of S<sup>2</sup>I-FSL mentioned in Section 1, we vary the sample ratio in  $\mathcal{U}$  as shown in Fig. 4. The accuracy decreases when fewer novel class samples are included in  $\mathcal{U}$ , since the more unbalanced samples between base and novel classes, the more challenging the task becomes. Compared to Meta-Inc-Baseline + PR, our method is more robust to changes in sample ratio, as less degradation ( $\delta$ ) is observed. More comparison and evaluation details are in Appendix G.




Figure 3: T-SNE [37] plots of queries and prototypes from a *tiered-ImageNet* test episode with and without our proposed components. Categories are represented by different colors.




Figure 4: Performance of different methods by varying the sample ratio in unlabeled set.

## 6 Conclusion

In this paper, a Semi-Supervised Incremental Few-Shot Learning (S<sup>2</sup>I-FSL) benchmark is proposed to generalize previous semi-supervised FSL and incremental FSL under a more realistic and challenging setting. To solve S<sup>2</sup>I-FSL, we introduce an efficient meta-baseline for incremental FSL. Based on that, a novel meta-training paradigm and a model adaptation scheme are proposed to fully explore unlabeled data. Extensive experiments on incremental FSL, semi-supervised FSL, and S<sup>2</sup>I-FSL demonstrate the effectiveness of our method.

## References

- [1] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. In *ICLR*, 2019.
- [2] Malik Boudiaf, Ziko Imtiaz Masud, Jérôme Rony, José Dolz, Pablo Piantanida, and Ismail Ben Ayed. Transductive information maximization for few-shot learning. In *NeurIPS*, 2020.
- [3] Kuilin Chen and Chi-Guhn Lee. Incremental few-shot learning via vector quantization in deep embedded space. In *ICLR*, 2021.
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [5] 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.
- [6] Yinbo Chen, Xiaolong Wang, Zhuang Liu, Huijuan Xu, and Trevor Darrell. A new meta-baseline for few-shot learning. *ArXiv*, abs/2003.04390, 2020.
- [7] Guneet S Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. In *ICLR*, 2020.
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [9] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018.
- [10] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Perez, and Matthieu Cord. Boosting few-shot visual learning with self-supervision. In *ICCV*, 2019.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [12] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [13] Yuqing Hu, Vincent Gripon, and Stéphane Pateux. Leveraging the feature distribution in transfer-based few-shot learning. *arXiv preprint arXiv:2006.03806*, 2020.
- [14] Huaxi Huang, Junjie Zhang, Jian Zhang, Qiang Wu, and Chang Xu. Ptn: A poisson transfer network for semi-supervised few-shot learning. In *AAAI*, 2021.
- [15] Anna Kukleva, Hilde Kuehne, and Bernt Schiele. Generalized and incremental few-shot learning by explicit learning and calibration without forgetting. In *ICCV*, 2021.
- [16] Kai Li, Yulun Zhang, Kunpeng Li, and Yun Fu. Adversarial feature hallucination networks for few-shot learning. In *CVPR*, 2020.
- [17] 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. *NeurIPS*, 2019.

- [18] Moshe Lichtenstein, Prasanna Sattigeri, Rogerio Feris, Raja Giryes, and Leonid Karlinsky. Tafssl: Task-adaptive feature sub-space learning for few-shot classification. In *ECCV*, 2020.
- [19] Jinlu Liu, Liang Song, and Yongqiang Qin. Prototype rectification for few-shot learning. In *ECCV*, 2020.
- [20] 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. In *ICLR*, 2019.
- [21] Pratik Mazumder, Pravendra Singh, and Piyush Rai. Few-shot lifelong learning. In *AAAI*, 2021.
- [22] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018.
- [23] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *CVPR*, 2018.
- [24] Limeng Qiao, Yemin Shi, Jia Li, Yaowei Wang, Tiejun Huang, and Yonghong Tian. Transductive episodic-wise adaptive metric for few-shot learning. In *ICCV*, 2019.
- [25] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L. Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, 2018.
- [26] Mengye Ren, Eleni Triantafyllou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018.
- [27] Mengye Ren, Renjie Liao, Ethan Fetaya, and Richard Zemel. Incremental few-shot learning with attention attractor networks. In *NeurIPS*, 2019.
- [28] Pau Rodríguez, Issam Laradji, Alexandre Drouin, and Alexandre Lacoste. Embedding propagation: Smoother manifold for few-shot classification. In *ECCV*, 2020.
- [29] 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.
- [30] Xiahan Shi, Leonard Salewski, Martin Schiegg, Zeynep Akata, and M. Welling. Relational generalized few-shot learning. In *BMVC*, 2020.
- [31] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [32] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020.
- [33] 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.

- [34] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. In *CVPR*, 2020.
- [35] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? In *ECCV*, 2020.
- [36] Hung-Yu Tseng, Hsin-Ying Lee, Jia-Bin Huang, and Ming-Hsuan Yang. Cross-domain few-shot classification via learned feature-wise transformation. In *ICLR*, 2020.
- [37] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [38] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, 2016.
- [39] Yikai Wang, Li Zhang, Yuan Yao, and Yanwei Fu. How to trust unlabeled data? instance credibility inference for few-shot learning. *arXiv preprint arXiv:2007.08461*, 2020.
- [40] Yu-Xiong Wang, Ross Girshick, Martial Hebert, and Bharath Hariharan. Low-shot learning from imaginary data. In *CVPR*, 2018.
- [41] 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.
- [42] Sung Whan Yoon, Jun Seo, and Jaekyun Moon. TapNet: Neural network augmented with task-adaptive projection for few-shot learning. In *ICML*, 2019.
- [43] Sung Whan Yoon, Do-Yeon Kim, Jun Seo, and Jaekyun Moon. Xtarnet: Learning to extract task-adaptive representation for incremental few-shot learning. In *ICML*, 2020.
- [44] Zhongjie Yu, Lin Chen, Zhongwei Cheng, and Jiebo Luo. Transmatch: A transfer-learning scheme for semi-supervised few-shot learning. In *CVPR*, 2020.
- [45] Baoquan Zhang, Xutao Li, Yunming Ye, Zhichao Huang, and Lisai Zhang. Prototype completion with primitive knowledge for few-shot learning. In *CVPR*, 2021.
- [46] Chi Zhang, Nan Song, Guosheng Lin, Yun Zheng, Pan Pan, and Yinghui Xu. Few-shot incremental learning with continually evolved classifiers. In *CVPR*, 2021.
- [47] Hongguang Zhang, Jing Zhang, and Piotr Koniusz. Few-shot learning via saliency-guided hallucination of samples. In *CVPR*, 2019.