

Incremental Few-Shot Object Detection via Simple Fine-Tuning Approach

Tae-Min Choi¹ and Jong-Hwan Kim^{1*}, *Fellow, IEEE*

Abstract—In this paper, we explore incremental few-shot object detection (iFSD), which incrementally learns novel classes using only a few examples without revisiting base classes. Previous iFSD works achieved the desired results by applying meta-learning. However, meta-learning approaches show insufficient performance that is difficult to apply to practical problems. In this light, we propose a simple fine-tuning-based approach, the Incremental Two-stage Fine-tuning Approach (iTFA) for iFSD, which contains three steps: 1) base training using abundant base classes with the class-agnostic box regressor, 2) separation of the RoI feature extractor and classifier into the base and novel class branches for preserving base knowledge, and 3) fine-tuning the novel branch using only a few novel class examples. We evaluate our iTFA on the real-world datasets PASCAL VOC, COCO, and LVIS. iTFA achieves competitive performance in COCO and shows a 30% higher AP accuracy than meta-learning methods in the LVIS dataset. Experimental results show the effectiveness and applicability of our proposed method¹.

I. INTRODUCTION

Deep learning using convolutional neural networks has recently made significant progress in computer vision [1], [2], [3], [4], [5]. These works, however, are usually developed in a supervised learning scheme that relies on abundant labeled datasets. Large-scale dataset-based training requires enormous human resources for data collection and annotation. It is also challenging to adapt models to novel classes with scarce data. As a result, few-shot learning has become a key area for many computer vision researchers.

In few-shot learning, it is difficult to train a model directly from a few novel data; therefore, we train general knowledge to the model with abundant base classes and transfer it to the novel classes. Yet, most of the work is focused on image classification tasks. Few-shot object detection (FSOD) is a more challenging task because it needs to train object recognition and localization using a small training set. Several methods [6], [7], [8], [9], [10], [11], [12], [13], [14] have been presented to overcome FSOD; meta-learning and fine-tuning based methods have been the main approaches to FSOD. Meta-learning studies have introduced a learning-to-learn scheme. They learn general knowledge by episodically training the model with base class data and transferring it to

the novel classes. These works can adapt quickly without extra training but have shown low performance that is difficult to apply in real-world applications. Alternatively, fine-tuning methods were developed based on the two-stage fine-tuning approach (TFA) [6] and outperformed meta-learning works.

Many FSOD methods usually only focus on the performance of the novel classes. Therefore, when adding the novel classes, they can not preserve the knowledge of the base classes. This leads to catastrophic forgetting. To overcome this issue, TFA fine-tunes the last layers of the object detector with a balanced dataset, including both the base and novel classes, to prevent forgetting the base classes. In reality, novel classes can be detected on the go due to the complex real-world environment. This creates problematic settings that cannot be revisited in the base classes. The FSOD methods degrade in these incremental learning settings, wherein they tend to forget base information when trained on a few novel data. In contrast, humans can learn a new concept incrementally using only small data and preserve prior knowledge without accessing the base data. Motivated by this gap between humans and machines, incremental few-shot object detection (iFSD) was proposed by [15]. iFSD uses the same dataset as FSOD with abundant base classes and a few novel classes. Also, the training strategy of iFSD is the same as FSOD; the model is first trained with base classes, and then the detector is incrementally learned using the novel classes. However, unlike FSOD, iFSD focuses on the performance of both the base and novel classes and transfers base knowledge to the novel classes without access to base class data.

In this paper, we propose a simple fine-tuning approach for iFSD, called iTFA. Inspired by TFA, we employ a two-stage training strategy based on Faster R-CNN [3]. We first train the Faster R-CNN with abundant base classes and then fine-tune the RoI feature extractor and classifier using a few novel classes. We focus on fine-tuning layer decisions and model design by analyzing the structure of Faster R-CNN and exploring the compatibility of the classifier optimized for iFSD.

TFA [6] prevented catastrophic forgetting due to a balanced fine-tuning training set, including the base classes. However, as mentioned above, we only access a few novel classes after training with the base classes in iFSD. Therefore, the two-stage training scheme must be modified to be suitable for iFSD. We train the object detector with a class-agnostic regressor using abundant base class data. The RoI feature extractor and classifier are separated and connected parallel to the base and novel branches. We also explore the effect of fine-tuning layer decision of the RoI feature

* Corresponding author

¹ Tae-Min Choi and Jong-Hwan Kim are with the School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon 34141, Korea (e-mail: tmchoi@rit.kaist.ac.kr; johkim@rit.kaist.ac.kr).

This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2020-0-00440, Development of Artificial Intelligence Technology that Continuously Improves Itself as the Situation Changes in the Real World)

¹Code is available at <https://github.com/TMIU/iTFA>

extractor and examine the compatibility of the classifier. With these modifications, iTFA outperforms the baselines on the novel class accuracy and overcomes catastrophic forgetting. In this paper, iTFA is evaluated on the PASCAL VOC [16], COCO [17], and LVIS [18] datasets for the iFSD task. Experiments show the effectiveness and competitive performance of our iTFA in all datasets.

To summarize, The contribution of this paper is three-fold: (1) We propose the Incremental Two-stage Fine-tuning Approach (iTFA), which is a simple fine-tuning strategy optimized for iFSD; (2) Concretely, we leverage a class-agnostic regressor trained on large-scale data. We also separate the base and novel branches to prevent catastrophic forgetting and empirically explore compatibility between the fine-tuning layer and classifier for the accurate performance on novel classes; (3) Extensive experiments on the iFSD benchmark demonstrate higher performance on novel classes and prevent degradation of the base class accuracy.

II. RELATED WORK

A. Few-shot object detection

FSOD aims to detect novel classes with a few labeled novel class samples and abundant base class samples. Many early works on FSOD apply meta-learning [13], [19], [12], [11], which uses the base samples to train a meta-learner to class agnostic knowledge so that they can be adapted to novel classes without extra training. Recently, [6] introduced TFA, which outperforms meta-learning-based methods. They presented a simple training strategy that fine-tunes only the last layers of the object detector. Following this approach, various fine-tuning-based methods [9], [8], [14], [7] have achieved state-of-the-art performance. Inspired by this workflow, we modify the TFA as a fine-tuning-based method for iFSD benchmarks.

B. Incremental few-shot object detection

iFSD was first proposed in [15]. Unlike FSOD, iFSD does not access the base class images when learning novel classes. Due to this assumption, catastrophic forgetting occurs; the model adds novel classes, then performance degradation occurs to the base classes. Like the FSOD workflow, prior works [15], [20] have developed meta-learning-based methods to solve these problems. ONCE [15] presented a class code generator to generate classification heads for novel classes. Sylph [20] developed a code predictor head and a code process module to solve iFSD without additional training and forgetting the base classes. Meta-learning-based methods can be directly adapted to novel classes; however, they have low accuracy. Better performance is needed for its application in a practical scenario. Fine-tuning-based works [21], [22] have been presented to achieve more accurate performance. They leverage knowledge distillation loss to overcome catastrophic forgetting. However, their network structure and training methods do not preserve the base class performance. To achieve higher accuracy in both the base and novel classes, we propose iTFA. This simple fine-tuning

approach is suitable for iFSD by modifying the structure and training strategy of TFA.

III. METHOD

A. Formulation

We follow the iFSD settings introduced in [6], [20]. In iFSD, there are two sets of dataset with base classes, C_{base} , and novel classes, C_{novel} . Base classes have abundant annotated training data, and novel classes have few training data (usually 1~10 per category). For an object detection dataset, $D = \{(x, y), x \in X, y \in Y\}$, where x is the input image and $y = \{(c_i, b_i), i = 1, \dots, N\}$ denotes the categories c and bounding box coordinates b of the N object instances in the image x . The goal of iFSD is to incrementally train the object detector to recognize both novel and base classes without revisiting base classes. As shown in Fig. 1, we split the training stage into the *base model training stage* and *fine-tuning stage*. We first pre-train the object detector with abundant base classes in the base model training stage. In this stage, we use D_{base} containing only base classes $c \in C_{base}$ with sufficient labeled images. After the first stage, we transfer prior knowledge to classify the novel classes that can use only K shots training data while maintaining the performance of base classes. In the fine-tuning stage, we use D_{novel} containing novel classes $c \in C_{novel}$ with K examples per class. In the test stage, our work aims to train a model that performs well on both base and novel classes, so we use D_{test} that contains all classes $c \in (C_{base} \cup C_{novel})$ at the test stage.

B. Incremental two-stage fine-tuning approach

In this section, we introduce our iTFA method for iFSD. As shown on the right side of Fig. 1(a), TFA uses base classes in the fine-tuning stage to prevent forgetting the prior knowledge. However, in iFSD, the object detector is fine-tuned only using novel classes without access to the base classes. TFA also freezes all structures before the box classifier, which disturbs the feature alignment of novel classes in the iFSD setting. To overcome these limitations, we propose iTFA, a simple fine-tuning approach optimized for iFSD. The training strategy of our iTFA has three steps: 1) changing the class-specific regressor to the class-agnostic regressor in the base training stage, 2) separating the RoI feature extractor (composed of two fc layers) and the classifier of Faster R-CNN [3] into the novel and base class branches, and 3) analyzing the effect of the decision of fine-tuning layers and compatibility of the classifier, then fine-tuning the novel branch.

1) *Base model training*: In the base training stage, we train the feature extractor, the classifier, and the box regressor using the base classes dataset D_{base} (left side of Fig. 1). We use the same training strategy from Faster R-CNN [3]. The loss term for training is as follows:

$$L = L_{rpm} + L_{cls} + L_{loc}, \quad (1)$$

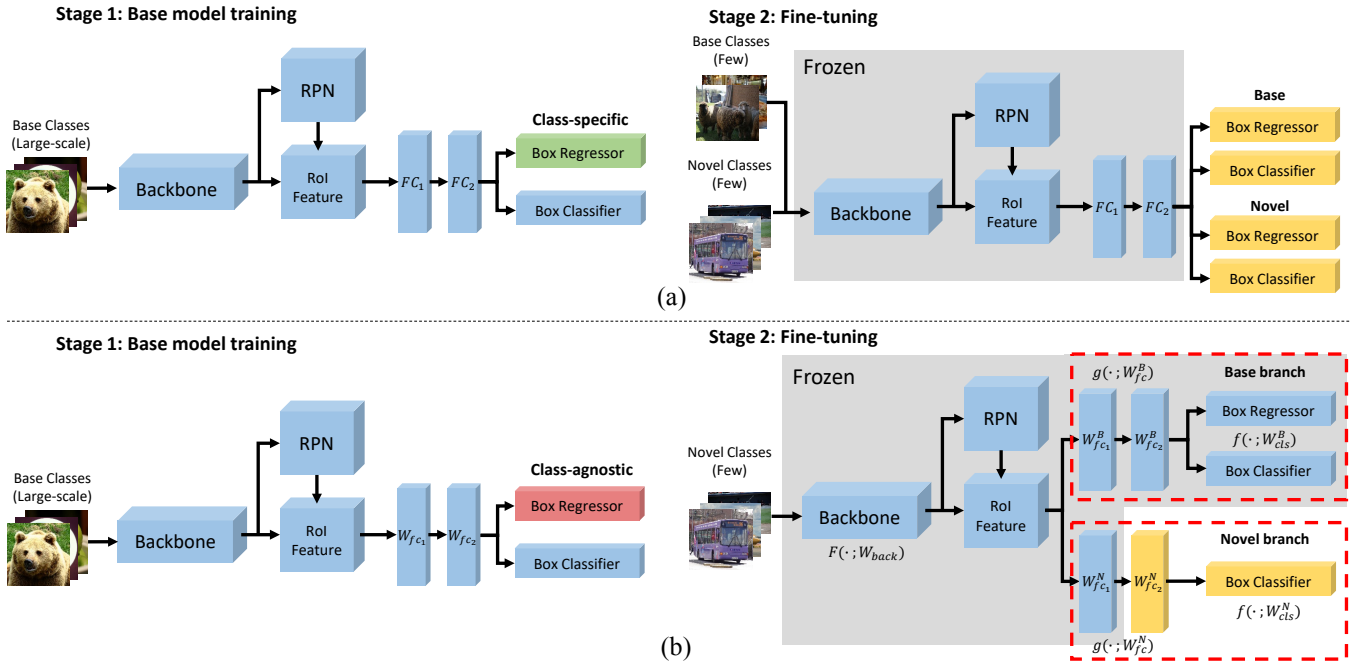


Fig. 1. Method overview. (a) the illustration of TFA, (b) the overall framework of iTFA. Both frameworks consist of two stages: base model training and fine-tuning stages. The whole object detector is trained with large-scale base classes in the base model training stage. In the fine-tuning stage, each method fixes the components indicated by gray boxes and fine-tunes other structures. TFA uses both base and novel classes in the fine-tuning stage. However, iTFA only uses novel classes, which causes catastrophic forgetting. To overcome catastrophic forgetting, iTFA divides the RoI feature extractor and classifier into two branches and fine-tunes only the novel branch. Also, the class-agnostic regressor is used for easier adaptation to novel classes.

where L_{rpn} is leveraged to determine whether the output of Region Proposal Network (RPN) is foreground or background and to adjust the anchor position, L_{cls} is the cross-entropy loss of the classifier, and L_{loc} is the L_1 loss for the box regressor.

2) *Class-agnostic box regressor*: Previous iFSD methods [15], [6], [22], [21] applied a class-specific box regressor to the object detector. A class-specific box regressor can generate a bounding box fitted to each class. However, it is challenging to train a class-specific regressor in few-shot learning due to the few instances of novel classes. It is difficult to obtain various bounding box information with a small number of instances because even the same class has different bounding box characteristics depending on the style and composition of the image. Therefore, we change the class-specific regressor to the class-agnostic box regressor in the base training stage and use it to predict the bounding boxes of both base and novel classes. The class-agnostic regressor can be leveraged for the novel classes because it has been pre-trained on the abundant dataset. When in the fine-tuning stage, the regressor is frozen, so it does not overfit to novel classes and keeps the knowledge of the base classes. Also, by freezing the class-agnostic box regressor, we can focus on classification rather than localization and prevent catastrophic forgetting of the base classes.

3) *Fine-tuning*: In the fine-tuning stage, the novel classes are trained incrementally with D_{novel} (K shots per class) without access to D_{base} (right side of Fig. 1(b)). Only the L_{cls} in (1) is used due to freezing the backbone, RPN, and class-agnostic regressor. In this stage, our goal is to maintain

the performance of base classes and adapt to the novel classes with a few images. In the following, we specify our fine-tuning strategy to achieve these goals.

4) *Separate novel branch*: TFA separates feature representation learning and box classifier learning into two stages. This means feature representations from the base classes could be transferred to novel classes without an extra weights update. The base model trained with sufficient base class data has a well-aligned feature space. Moreover, the classifier, which is fine-tuned with a few shot data including both C_{base} and C_{novel} , classifies novel classes without forgetting the decision boundary of the base classes. However, in iFSD, we cannot use C_{base} during fine-tuning. If the base classifier is fine-tuned without base class data, it will cause overfitting on the novel classes and lead to catastrophic forgetting of the base classes. Therefore, we first separate the base and novel classifiers and then freeze the base classifier and fine-tune only the novel classifier. Separating the novel and base classifiers allows us to focus on feature alignment for novel classes only. For this purpose, as shown in Fig. 1(b), we add the RoI feature extractor (W_{fc1}^N and W_{fc2}^N) for the novel classes in parallel, and it is trained to align the novel class features. We only separate the RoI features generated from the backbone and RPN have sufficient information about the objects in the image due to being trained with abundant base classes.

We design the fine-tuning strategy to align the novel

feature space as follows:

$$\min_{W_{fc_2}^N, W_{cls}^N} L_{cls}(y_i, [p_i^B, p_i^N]), \quad (2)$$

$$\text{with } p_i^B = f(g(z_i; W_{fc}^B); W_{cls}^B),$$

$$p_i^N = f(g(z_i; W_{fc}^N); W_{cls}^N), z_i = \mathcal{F}(x_i; W_{back}),$$

where $f(\cdot; W_{cls}^B)$ and $f(\cdot; W_{cls}^N)$ are the base and novel classifiers, respectively. $g(\cdot; W_{fc}^B)$ and $g(\cdot; W_{fc}^N)$ are the RoI feature extractors for the base and novel classes, respectively. $W_{fc_2}^N$ is the last fc layer of the novel RoI feature extractor. W_{back} is the backbone network, and z_i means representation vector created from the backbone. p_i^B is $(|C_{base}|+1)$ -way including the background as base classes prediction, and p_i^N is $|C_{novel}|$ -way as novel classes prediction. Then, L_{cls} is calculated by concatenating the two predictions. W_{cls}^N is randomly initialized and pre-trained weights W_{fc_1} and W_{fc_2} from the base training stage are loaded into $W_{fc_1}^N$ and $W_{fc_2}^N$, respectively. Other networks are frozen except W_{cls}^N and $W_{fc_2}^N$ to preserve their prior knowledge. In particular, as shown in Fig. 1(b), we fine-tune only $W_{fc_2}^N$ in the fine-tuning stage because, empirically, fine-tuning $W_{fc_1}^N$ is found to disturb the disentangling of novel features. The effect of fine-tuning layers is evaluated in Section IV-B.

5) *Compatibility of classifier*: Since the feature spaces of the base and novel classes are separated, we consider the compatibility of each RoI feature and classifier. In few-shot learning, there is a bias in which classifiers predict base classes with higher scores, which makes it more challenging to classify novel classes. Therefore, prior few-shot learning works use cosine classifiers [23], [24], [25] with normalized features and various margin losses [26], [27] to align intra-class and inter-class distributions of the novel classes. Cosine similarity-based classifiers are effective for FSOD since they focus on disentangling only the feature space of the novel class. However, in iFSD, we need to align the novel class features without disturbing the feature space of the base classes. In our method, making enough space for the novel classes on the hypersphere is difficult due to freezing the base class feature space during the fine-tuning stage. The inter-class separability is violated because there is no realignment on the base class feature space. Thus, we test the compatibility of the classifier and empirically find that a linear classifier helps disentangle the novel class feature and improve the accuracy of novel classes.

IV. EXPERIMENTS

A. Implementation details

1) *Datasets and metric*: We evaluated iTFA on PASCAL VOC 2012 [16], COCO [17], and LVIS v1 [18] datasets. The evaluation metric and the data split strategy were followed from [6], [15], [20] for a fair comparison. For COCO, 60 categories disjoint from PASCAL VOC are considered base classes; the other 20 classes are considered novel classes. For LVIS v1, 1203 categories are composed of 337 rare, 461 common, and 405 frequent classes. The 337 rare categories

were treated as novel classes, and the 866 common and frequent categories as base classes. For the cross-dataset evaluation from COCO to PASCAL VOC, we used the same base and novel class set (60 and 20 categories from COCO) for the base model training and fine-tuning stage, respectively. We then evaluated the object detector on the PASCAL VOC 2012 test set. We evaluated our work on 1, 5, and 10 shots and used AP50 and the standard COCO-style metric, which calculates the average precision (AP) at multiple intersection-over-union (IoU) thresholds from 0.5 to 0.95. For stable results, ten tests were made on each random seed, and the mean average precision was reported for all experiments. For COCO, base, novel, and harmonic mean AP (mAP) were reported as bAP, nAP, and hAP, respectively. For LVIS, APr, APc, and APf are the mAP of rare, common, and frequent categories, respectively. Also, APs, APm, and API are the mAP of small, medium, and large objects. For cross-dataset evaluation, AP50 was used for comparison to the other methods.

2) *Training setting*: Our method was implemented using MMDetection [28]. We used Faster R-CNN [3] with Feature Pyramid Network [29] for all experiments. The ResNet-50/101 [30] backbones were used for COCO, while ResNet-50 and ResNet-101 were used for LVIS and Pascal VOC for the fair comparison, respectively. SGD with a momentum of 0.9, a weight decay of $1e-4$, and a batch size of 8 was used on 4 NVIDIA RTX 2080Ti GPUs, with two images per GPU. We set the learning rate to $1e-2$ for both the base training and the fine-tuning stages. For COCO and PASCAL VOC, we trained the base model for 12 epochs and fine-tuned the object detector for 3000 steps. The learning rate decreased ten-fold at epochs 8 and 11 in the base training and at steps 1500 and 2000 in the fine-tuning stage. For LVIS, we trained the base model for 24 epochs, and the learning rate decreased at epochs 16 and 22. We fine-tuned the model for 10000 iterations, and the learning rate decreased at 7000 and 9000 iterations.

B. Ablation study

We evaluated the following ablations on COCO to analyze how each component of iTFA affects the novel class performance. All training settings were maintained as mentioned in Section IV-A. The performance was evaluated by changing three components: the fine-tuning layer selection, regressor, and classifier. The fine-tuning layer decides which layers of the RoI feature extractor to fine-tune. *None* means that all layers are frozen, and only the classifier is updated. fc_1 and fc_2 indicate the fine-tuning $W_{fc_1}^N$ and $W_{fc_2}^N$, respectively. In addition, we tested the effects of class-specific and class-agnostic regressors. The effect of the classifier was tested using the linear and cosine classifier. The scale factor of the cosine classifier was set to 20.

As Table I shows, fine-tuning only $W_{fc_2}^N$ with a linear classifier and class-agnostic regressor outperforms other settings in all shots. This demonstrates that fine-tuning $W_{fc_2}^N$ with freezing $W_{fc_1}^N$ helps disentangle the novel class features. Also, a cosine classifier with normalized features disturbs

TABLE I

ABLATION STUDY OF FINE-TUNING LAYER DECISION, CLASSIFIER, AND REGRESSOR OF iTFA. **BOLD** AND UNDERLINE INDICATE THE BEST AND THE SECOND BEST PERFORMANCE, RESPECTIVELY.

Fine-tuning layer			Regressor		Classifier		Shots		
<i>None</i>	fc_1	fc_2	specific	agnostic	Linear	Cosine	1	5	10
✓				✓	✓		3.4	7.0	8.3
		✓		✓	✓		4.3	9.9	11.8
	✓	✓		✓	✓		<u>3.6</u>	<u>9.0</u>	10.9
✓				✓		✓	<u>3.6</u>	7.6	9.2
		✓		✓		✓	3.3	8.9	<u>11.0</u>
	✓	✓		✓		✓	2.6	8.4	10.5
		✓	✓		✓		2.8	7.5	9.7
		✓	✓			✓	3.0	8.2	10.2

TABLE II

PERFORMANCE OF NOVEL AND BASE CLASSES ON THE COCO DATASET.

Shots	Method	Backbone	Novel		Base		Harmonic	
			nAP	nAP50	bAP	bAP50	hAP	hAP50
1	ONCE [15]	ResNet-50	0.7	-	17.9	-	1.4	-
	Sylph [20]		1.1	-	37.6	-	2.1	-
	Ours		3.8	6.5	35.7	55.9	6.8	11.7
	TFA [6]	ResNet-101	1.9	3.8	31.9	51.8	3.6	7.1
	iMTFA [31]		3.2	5.9	27.8	40.1	5.8	10.3
	LEAST [21]		<u>4.2</u>	-	29.5	-	<u>7.4</u>	-
	Ours		4.3	7.2	<u>37.4</u>	57.4	7.7	12.8
5	ONCE [15]	ResNet-50	1.0	-	17.9	-	1.9	-
	Sylph [20]		1.5	-	42.4	-	2.9	-
	Ours		8.3	15.7	35.5	56.0	13.5	24.5
	TFA [6]	ResNet-101	7.0	13.3	32.3	50.5	11.5	21.1
	iMTFA [31]		6.1	11.2	24.1	33.7	9.7	16.8
	LEAST [21]		<u>9.3</u>	-	31.3	-	<u>14.3</u>	-
	Ours		9.9	17.8	<u>37.2</u>	57.3	15.6	27.2
10	ONCE [15]	ResNet-50	1.2	-	17.9	-	2.3	-
	Sylph [20]		1.7	-	42.8	-	3.3	-
	Ours		10.2	19.3	35.5	56.0	15.9	28.7
	TFA [6]	ResNet-101	9.1	17.1	32.4	50.6	14.2	25.6
	iMTFA [31]		7.0	12.7	23.4	32.4	10.7	18.3
	DBF [22]		9.1	-	28.5	-	13.8	-
	LEAST [21]		12.8	-	31.3	-	18.2	-
Ours	<u>11.8</u>	21.3	<u>37.2</u>	57.3	<u>17.9</u>	31.1		

the novel feature alignment. Compared to rows 2, 5, and the last two rows of Table I, Faster R-CNN with a class-agnostic regressor outperforms a class-specific regressor in all shots. This justifies that the class-agnostic regressor can be leveraged to the novel classes and is more suitable for iFSD tasks.

C. Comparison with prior works

1) *Baselines*: We considered six baselines in our experiments. TFA [6] is a fine-tuning-based approach for FSOD. ONCE [15] and Sylph [20] are the meta-learning-based approaches for iFSD; they use a one-stage object detector due to its superior generalization performance. iMTFA [31] is designed for incremental few-shot semantic segmentation based on Mask R-CNN [4]; it creates novel classifier weights through weight imprinting. DBF [22], and LEAST [21] use

Faster R-CNN as fine-tuning-based iFSD methods. DBF presents a progressive model updating rule to preserve the long-term memory of old classes. LEAST stores a few exemplars from the base training set to preserve prior knowledge when fine-tuning the detector.

2) *Results on COCO*: As shown in Table II, our iTFA achieves competitive performance in all settings. In particular, iTFA outperforms TFA by 20-40% on 5 and 10-shots and even doubles the *nAP* on a 1-shot setting. This demonstrates that iTFA is a more suitable fine-tuning approach for iFSD tasks. In addition, our method surpasses meta-learning methods by a large margin in all environments. This corroborates that our training strategy can be applied to practical iFSD tasks. Our method also shows consistent base class performance in all shots, demonstrating that catastrophic forgetting is prevented by disentangling the feature

space of the base and novel classes by separating the RoI feature extractor and classifier. Sylph shows the highest bAP in all settings. This performance difference is due to the type of object detector. Sylph used the FCOS [32] as the base detector, which outperforms Faster R-CNN. In this paper, our objective is to propose a fine-tuning strategy for the iFSD task, so we used Faster R-CNN, the most popular two-stage object detector, as the base detector. In addition, although we did not use the base class exemplars in the fine-tuning stage, we achieved more accurate performance in the 1-shot and 5-shot settings and competitive performance compared to LEAST in the 10-shot setting. LEAST shows a higher nAP at the 10-shot. LEAST stores k base classes exemplars (in k -shot) and uses them for fine-tuning to keep the intra-class variance of base classes. It helps distinguish novel and base classes in the fine-tuning stage, so LEAST performs better in a 10-shot setting that can store many base exemplars. However, saving the base class exemplars is not a fair comparison because it violates the iFSD assumption. Despite using a basic detector and only novel class data for fine-tuning, iTFA shows almost the best performance for several environments. These results demonstrate the strength of our iTFA in preserving knowledge and accurate adaptation with a few samples.

3) *Results on LVIS v1*: Our iTFA was also evaluated on the LVIS dataset. LVIS is a natural long-tail distribution dataset with a large number of categories. Since rare classes of LVIS have 1 to 10 instances, each class is sampled uniformly with 10 instances to balance the novel set. Table III shows our results compared to recent works on LVIS. We tested iTFA on the iFSD setting that treated rare classes ($K < 10$) to a novel set. Our method significantly outperforms meta-learning methods with about +4 and +12 AP gains on rare classes. Although iTFA uses only novel classes in the fine-tuning stage, it showed competitive performance in all categories against TFA and the joint-train model. This demonstrates the effectiveness of iTFA on the long-tail distribution dataset.

TABLE III

PERFORMANCE OF RARE, COMMON, AND FREQUENT CLASSES ON THE LVIS v1 DATASET.

Method	AP	APs	APm	APl	APr	APc	APf
Joint-train [18]	23.5	18.2	30.9	35.4	10.5	<u>22.2</u>	30.6
TFA [6]	24.4	19.9	29.5	38.2	<u>16.9</u>	24.3	27.7
ONCE [15]	12.9	-	-	-	6.3	11.2	17.7
Sylph [20]	20.7	-	-	-	13.9	19.0	25.5
Ours	<u>24.0</u>	17.5	31.2	35.9	18.1	21.0	<u>29.9</u>

4) *Results on cross-domain from COCO to PASCAL VOC*: We evaluated iTFA in the cross-dataset setting from COCO to PASCAL VOC. Since PASCAL VOC has no base classes, we only report the novel class performance. Table IV shows the same relative results evaluated on COCO. Our method performs well against the prior methods in all the settings, indicating the strength of our method.

5) *Qualitative results*: Fig. 2 illustrates the qualitative results of our method on the COCO dataset. The top two

rows show success results, and the bottom row shows failure results. The success cases show that base (blue boxes) and novel classes (red boxes) are correctly detected in one image. In the bottom row, there are some misclassified results for failure cases due to similar characteristics and appearance. From left to right: the boat is misclassified as an airplane, the birds are confused for sheep and cow, and the train is detected well, but the bus is detected even though there is no bus.

TABLE IV

PERFORMANCE OF THE CROSS-DOMAIN EXPERIMENT FROM COCO TO PASCAL VOC.

Shots	Method				
	Ours	DBF [22]	LEAST [21]	LEAST-NE [21]	ONCE [15]
1	11.3	-	6.8	<u>7.2</u>	-
5	26.0	12.5	<u>16.7</u>	16.5	2.4
10	30.4	<u>23.7</u>	18.8	18.5	2.6



Fig. 2. Qualitative results of iTFA on COCO with $K=10$. The top two rows are success cases, and the bottom row images are failure cases. The blue and red boxes indicate the base and novel classes, respectively (zoom view better).

V. CONCLUSIONS

In this paper, we proposed iTFA for iFSD. To solve iFSD, we addressed a simple fine-tuning strategy by modifying TFA. We divided the RoI feature extractor and classifier into the base and novel branches and then fine-tuned only the novel branch. We also adopted a class-agnostic regressor and analyzed the compatibility of the classifier. Our method outperformed the prior meta-learning methods by a large margin in widely used datasets. Although our method fine-tuned the model without exemplars of the base classes, it showed competitive results with other fine-tuning-based methods. We will adapt our fine-tuning strategy to various object detector structures for further improvement.

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, 2012.
- [2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [3] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, 2015.
- [4] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.
- [5] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [6] X. Wang, T. Huang, J. Gonzalez, T. Darrell, and F. Yu, "Frustratingly simple few-shot object detection," in *International Conference on Machine Learning*. PMLR, 2020, pp. 9919–9928.
- [7] C. Xing, N. Rostamzadeh, B. Oreshkin, and P. O. O. Pinheiro, "Adaptive cross-modal few-shot learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [8] J. Wu, S. Liu, D. Huang, and Y. Wang, "Multi-scale positive sample refinement for few-shot object detection," in *European conference on computer vision*. Springer, 2020, pp. 456–472.
- [9] B. Sun, B. Li, S. Cai, Y. Yuan, and C. Zhang, "Fsce: Few-shot object detection via contrastive proposal encoding," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7352–7362.
- [10] J. Kim, I. Yoon, G.-M. Park, and J.-H. Kim, "Non-probabilistic cosine similarity loss for few-shot image classification," in *BMVC*, 2020.
- [11] L. Karlinsky, J. Shtok, S. Harary, E. Schwartz, A. Aides, R. Feris, R. Giryes, and A. M. Bronstein, "Repmet: Representative-based metric learning for classification and few-shot object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5197–5206.
- [12] Y. Xiao and R. Marlet, "Few-shot object detection and viewpoint estimation for objects in the wild," in *European conference on computer vision*. Springer, 2020, pp. 192–210.
- [13] B. Kang, Z. Liu, X. Wang, F. Yu, J. Feng, and T. Darrell, "Few-shot object detection via feature reweighting," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 8420–8429.
- [14] Y. Cao, J. Wang, Y. Jin, T. Wu, K. Chen, Z. Liu, and D. Lin, "Few-shot object detection via association and discrimination," *Advances in Neural Information Processing Systems*, vol. 34, pp. 16570–16581, 2021.
- [15] J.-M. Perez-Rua, X. Zhu, T. M. Hospedales, and T. Xiang, "Incremental few-shot object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 13 846–13 855.
- [16] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [17] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.
- [18] A. Gupta, P. Dollar, and R. Girshick, "Lvis: A dataset for large vocabulary instance segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5356–5364.
- [19] X. Yan, Z. Chen, A. Xu, X. Wang, X. Liang, and L. Lin, "Meta r-cnn: Towards general solver for instance-level low-shot learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9577–9586.
- [20] L. Yin, J. M. Perez-Rua, and K. J. Liang, "Sylph: A hypernetwork framework for incremental few-shot object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9035–9045.
- [21] P. Li, Y. Li, H. Cui, and D. Wang, "Class-incremental few-shot object detection," *arXiv preprint arXiv:2105.07637*, 2021.
- [22] Y. Li, H. Zhu, J. Ma, C. S. Teo, C. Xiang, P. Vadakkepat, and T. H. Lee, "Towards generalized and incremental few-shot object detection," *arXiv preprint arXiv:2109.11336*, 2021.
- [23] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4367–4375.
- [24] H. Qi, M. Brown, and D. G. Lowe, "Low-shot learning with imprinted weights," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 5822–5830.
- [25] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang, "A closer look at few-shot classification," *arXiv preprint arXiv:1904.04232*, 2019.
- [26] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4690–4699.
- [27] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "Cosface: Large margin cosine loss for deep face recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 5265–5274.
- [28] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin, "MMDetection: Open mmlab detection toolbox and benchmark," *arXiv preprint arXiv:1906.07155*, 2019.
- [29] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2117–2125.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [31] D. A. Ganea, B. Boom, and R. Poppe, "Incremental few-shot instance segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1185–1194.
- [32] Z. Tian, C. Shen, H. Chen, and T. He, "Fcoss: Fully convolutional one-stage object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9627–9636.