

# Sylph: A Hypernetwork Framework for Incremental Few-shot Object Detection

Li Yin Juan M Perez-Rua Kevin J Liang  
Meta AI

{liyin, jmp, kevinjliang}@fb.com

## Abstract

We study the challenging incremental few-shot object detection (iFSD) setting. Recently, hypernetwork-based approaches have been studied in the context of continuous and finetune-free iFSD with limited success. We take a closer look at important design choices of such methods, leading to several key improvements and resulting in a more accurate and flexible framework, which we call Sylph. In particular, we demonstrate the effectiveness of decoupling object classification from localization by leveraging a base detector that is pretrained for class-agnostic localization on large-scale dataset. Contrary to what previous results have suggested, we show that with a carefully designed class-conditional hypernetwork, finetune-free iFSD can be highly effective, especially when a large number of base categories with abundant data are available for meta-training, almost approaching alternatives that undergo test-time-training. This result is even more significant considering its many practical advantages: (1) incrementally learning new classes in sequence without additional training, (2) detecting both novel and seen classes in a single pass, and (3) no forgetting of previously seen classes. We benchmark our model on both COCO and LVIS, reporting as high as 17% AP on the long-tail rare classes on LVIS, indicating the promise of hypernetwork-based iFSD.

## 1. Introduction

While advances in deep learning have led to significant progress in computer vision [18, 23, 24, 31], much of this success has relied upon large-scale data collection and annotation [7, 20, 22, 36], a process that is both labor-intensive and time-consuming, and does not scale well with the number of categories. This is especially true for object detection [18, 23, 37], particularly for the long tail of object categories, where data may be scarcer [22]. As a result, few-shot learning of object detectors (FSD) [28, 65, 70, 72] has become a recent topic of interest.

While learning a novel class from only a few samples

alone is a challenging problem, the task can be made simpler by leveraging known classes with abundant data (commonly referred to as *base* classes), whose structure can be used as a prior for knowledge transfer. The few previous FSD works have approached this primarily in two ways. The first is fine-tuning [47, 65], where a model is first pretrained on the base classes and then fine-tuned on a small balanced set of data from both the base and novel classes, a form of test-time training [59]. Although simple, it has difficulty scaling to many real-world applications due to its computational and memory requirements. An alternate strategy is taking a meta-learning approach [72]. Meta-learning approaches frame the problem as “learning to learn” [4, 10, 32, 44, 61, 69, 72], training the model episodically to induce fast adaptation to novel classes.

However, many FSD methods focus on the limited set-up where only novel categories are to be detected. These methods often fail to preserve the original detector performance on base categories [4, 10, 32, 72] or forget about the ones it was initially trained on [65]. Given the ever-changing nature of the real-world, a desirable property of machine learning systems is the ability to incrementally learn new concepts without revisiting previous ones and not forgetting them [40, 42]. Humans are able to achieve such feat, learning novel concepts not only without forgetting but reusing such knowledge [45]. Conventional supervised learning struggles with incrementally presented data, tending to suffer catastrophic forgetting [39, 51]. An alternative is studying all the available data every time new concepts arrive, commonly referred to as “joint training” [22], but such a paradigm imposes a slow development cycle, requiring significant data collection efforts for the new concepts and expensive large-scale training (and re-training).

Instead, we seek an object detection model capable of learning new classes from a few shots in a fast, scalable manner without forgetting previously seen classes, a setting commonly referred to as incremental few-shot detection (iFSD). ONCE [44], a meta-learning approach to FSD, is of particular interest due to its hypernetwork-based class-conditional design. ONCE is able to enroll novel categories

without affecting its ability to remember base classes. We use a base detector and hypernetwork architecture similar to ONCE, but with a few key design differences: (1) ONCE, along several other recent works [28, 65, 72], attempts to directly produce (via training or hypernet) the parameters of a localization regression model that transforms the query sample feature maps into the output bounding boxes, all from the few available training samples. We find this to be unnecessary and potentially harmful, as the task can be significantly simplified by decoupling localization from classification. To achieve this goal, we leverage a base detector with class-agnostic localization capability pretrained on abundant base class data. (2) We study the class-conditional hypernetwork’s behavior, making some key changes to the structure and adding normalization to the predicted parameters, resulting in much higher accuracy.

With an architecture that can swiftly adapt to the long tail of classes from few shots, we name our framework *Sylph*, after the nimble long-tailed hummingbird (Figure 1). We present extensive evaluations that empirically demonstrate the benefits of our design, showing that Sylph is more effective than ONCE [44] (our main baseline) across all the reported datasets and evaluation regimes. On the challenging LVIS few-shot learning benchmark in particular, we show that Sylph is superior by a margin of 8% points.

## 2. Related Work

**Object Detection** Object detection is the task of simultaneously localizing and classifying objects within a scene. Most modern object detectors consist of a convolutional feature extractor [24, 31, 57] followed by various mechanisms or networks to predict classes and some form of bounding box coordinates [25]. Detectors that first generate region proposals during inference are often referred to as two-stage detectors [6, 17, 18, 23, 54], while ones that directly predict class and localization from the convolutional feature maps are considered single-stage detectors [8, 35, 38, 53, 62]. Single-stage detectors have the advantage of having simpler implementations and faster inference speeds, and recent advances have increased their accuracy to be competitive with two-stage models [37], which had previously been the primary advantage of such models. Throughout this work, we choose to primarily use FCOS [62] as our base detector due to its strong performance and class-agnostic localization based on “centerness” and intersection-over-union (IoU) losses; this allows for better generalization and high recall on novel unseen classes [29], especially when trained on large-scale datasets.

**Few-shot Learning** While many supervised learning approaches assume a large number of samples from the data distribution, such methods risk overfitting when the model has only a few samples to learn from. Given the costs of collecting, annotating, and training models with large amounts

of data, few-shot learning has become an active research direction, with image classification as the most common task. Many recent approaches take a meta-learning strategy [63]. Optimization-based approaches produce models that can quickly learn from few samples [13, 41, 52]. Metric-learning methods learn an embedding function that induces a space where samples can be compared with nearest neighbors or other such simple algorithms [33, 58, 60, 64]. Hypernetworks have also been used to predict model parameters for new classes from limited samples [3, 14, 16, 48, 50, 66]. We use a hypernetwork in our model to predict convolutional kernels for novel object classification. Such a strategy requires zero training during inference time and can easily scale to an arbitrary number of classes.

**Few-shot Object Detection and Beyond** Most neural networks are trained with stochastic gradient descent, which often assumes the training data are drawn independently and identically distributed (*i.i.d.*). However, this *i.i.d.* assumption is violated in the practical scenario where few-shot categories are seen only after the model have been trained for a set of base categories. In such situations, catastrophic forgetting [12, 19] can occur: the model suffers severe degradation in performance on the original classes. In image classification, some works have proposed a generalized setting for few-shot learning to tackle this exact situation [16, 46]. Similarly for object detection, recent works have focused on incorporating few-shot categories into a model that has been pretrained with large-scale datasets [11, 44, 65]. This goes beyond the simpler more traditional few-shot object detection set-up [28, 47, 48]. More generally, continual object detection [1, 44] works attempt to learn to detect new classes through several learning instances without forgetting any of the seen categories.

Of the prior work with the goal of both few-shot and continual learning for object detection, some are continual only in that they do not degrade base class accuracy during a *single* few-shot adaptation to new classes [11, 65]. In contrast, ONCE [44] considers a setting in which novel classes arrive sequentially and *incrementally*, leading to multiple learning events during which forgetting must be avoided. We adopt a model architecture that is able to provide such capabilities, as it is more flexible and a better fit for interactively learning new classes from the world. Methodologically, however, we approach the problem differently, as we (1) simplify learning by utilizing a detector with class-agnostic localization rather than trying to learn per-class localization from only a few samples; (2) leverage a per-class binary classifier to allow incrementally and independently added novel classes to co-exist with previously learned base classes, detecting seen and novel classes in a single pass; (3) generate both weights and biases for newly added classes, proposing an effective weight normalization to the output of a hypernetwork weight generator that enables stable training and more

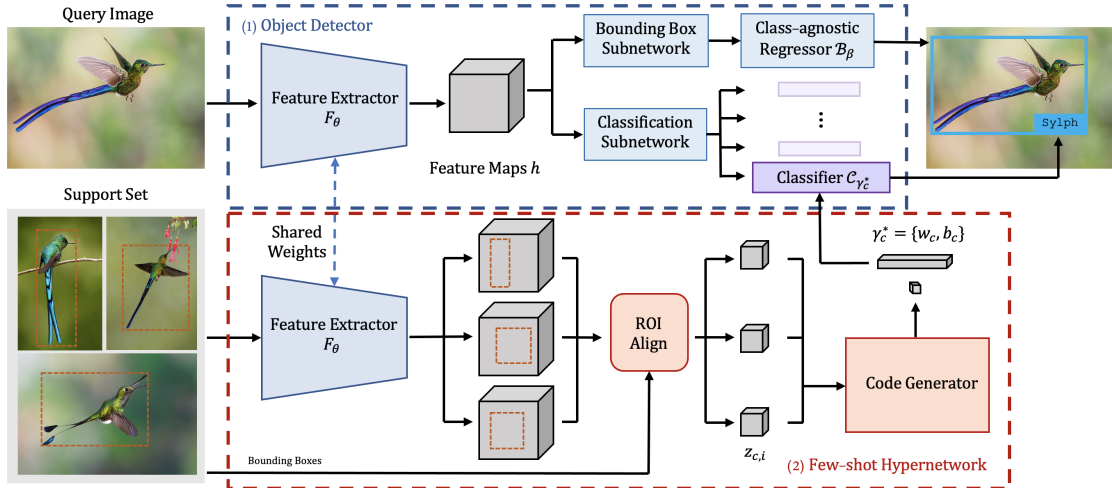


Figure 1. **The Sylph Framework.** Sylph is composed of a base object detector and a few-shot hypernetwork, whose Code Generator consists of a Code Predictor Head and Code Process Module (detailed in Section. 3.1.2). The dashed arrow indicates weight sharing.

effective synthesis of class-specific class codes.

### 3. Methods

We seek a model that can operate in the incremental few-shot detection (iFSD) [44] setting: a detector that can flexibly adapt to new classes introduced in sequence from only a few examples, without forgetting any previously seen classes. We differentiate this *continuous* iFSD with *batch* iFSD where novel classes are added in a batch. Concretely, after being pretrained on a base set of classes  $C^b$ , the objective is to achieve good performance on a novel class  $c_t^n \in C^n$  from a support set of only  $K$  shots while maintaining strong performance on  $C^b$  and the preceding novel classes  $c_{t'}^n \forall t' < t$ , *without* re-training on data from these previous classes. As the goal is to learn to adapt to new classes, we assume  $C^b \cap C^n = \emptyset$ .

#### 3.1. Sylph

To achieve the stated objective of iFSD, we introduce Sylph, a framework that can quickly add detection capabilities of new classes, without any additional optimization of model parameters. Sylph is composed of two primary components (Figure 1): (1) a base object detector with class-agnostic localization to surface salient objects in an image with high recall and (2) a few-shot hypernetwork to generate class-specific parameters for a per-class binary classifier. We discuss each of these in detail below.

##### 3.1.1 Object Detector

Modern object detection models [25] are often composed of a convolutional backbone  $\mathcal{F}_\theta$  followed by a detector head  $\mathcal{D}_\phi$ . Given an image  $I$ , the former produces high-level fea-

ture maps  $h = \mathcal{F}_\theta(I)$ , which can then be used by the detector head to predict both class  $c$  and location, as specified by a bounding box  $b = (x, y, h, w)$ . Many detection models perform both these tasks in parallel [18, 38, 54], predicting the class category and bounding box coordinates from the same features:  $\mathbf{o} = \mathcal{D}_\phi(h)$ , where  $\mathbf{o} = [o_1, \dots, o_n]$  are predicted objects in  $I$ , with each object  $o_i = [c_i, b_i]$  containing the predicted class label and bounding box. We denote the final regression and classification layer as  $\mathcal{B}_\beta$  and  $\mathcal{C}_\gamma$ , which can be a fully connected layer in region-based detection [54] or a convolutional layer in dense prediction [62]. For an  $N$ -way classification problem, the parameters  $\gamma$  for the classifier normally produce  $N + 1$  logits for a softmax, corresponding to the  $N$  classes and the background. Meanwhile, the bounding box regressor’s parameters  $\beta$  contains  $N$  stacked weights  $\beta_c$ , with one for each class  $c$ ; the class with the highest prediction score determines which regressor’s prediction is selected. In order for our object detector to support the challenging iFSD setting, we make several key design choices affecting the two primary outputs of a detector: classification and localization.

**Incremental Classification Without Forgetting** A major contributor to catastrophic forgetting is highly non-*i.i.d.* sequential training with a shared classification head [12]; optimizing the softmax can result in destructive gradients overwriting previous knowledge. We thus replace the single softmax-based classifier  $\mathcal{C}_\gamma$  with many binary sigmoid-based classifiers  $\mathcal{C}_{\gamma_c}$ , with each class individually handled by its own set of parameters. When trained with the focal loss [35], sigmoid classifiers have been shown to be just as effective as a single softmax classifier. Thus, when adding novel classes, we can train or generate a new set of classifier parameters  $\gamma_c^n$ . When combined with previous parameters

to predict all available classes, there is zero interference between each class’s prediction score.

**Class-agnostic Bounding Box Regressor** Previous few-shot object detection methods [28, 44, 65] have tended to learn a per-class box regressor  $\mathcal{B}_{\beta_c}$  in tandem with the classifier. However, when only a few examples are available for learning, the model has very little opportunity to learn a custom location regressor for each novel class. Instead, we propose pretraining the base object detector with a single class-agnostic box regressor  $\mathcal{B}_{\beta}$  for all classes. When adapting the model to novel classes  $C^n$ , we simply reuse  $\mathcal{B}_{\beta}$  for localization. Such an approach has been shown to work well for zero-shot object detection if pretrained on a large-scale dataset [21] and can leverage progress in the open-world detection literature [29]. By alleviating the need to learn localization in a few-shot or continual manner, we can treat the problem as a few-shot classification task and focus just on generating additional classifier parameters  $\gamma_c^n$ . We validate the effectiveness of this setup in Section 5.

We can satisfy both the aforementioned objectives with FCOS [62], a simple one-stage and anchor-free object detector. With these design choices, we decouple the few-shot novel class detection problem into serial tasks of localization and few-shot classification, dramatically simplifying it.

### 3.1.2 Few-shot Hypernetwork

With localization handled by the class-agnostic object detector, the problem reduces to few-shot classification. Sylph uses a hypernetwork  $\mathcal{H}_{\psi}$  to generate parameters  $\gamma_c^* = \{w_c, b_c\}$  for each binary classifier  $\mathcal{C}_{\gamma_c^*}$ .  $\mathcal{H}_{\psi}$  takes as input an  $N$ -way  $K$ -shot episode of support set samples, consisting of  $K$  instances of  $N$  classes randomly sampled from the meta-training set. We denote this support set  $S^{N \times K} = (I^{N \times K}, b^{N \times K})$ , with  $I^{N \times K} \in \mathbb{R}^{(N \times K) \times C \times H \times W}$ . The hypernetwork is modularized into three components: support set feature extraction, code prediction, and code aggregation and normalization, which we detail below.

**Support Set Feature Extraction** The first stage consists of extracting features from the episode’s support set. We share the same convolutional backbone  $\mathcal{F}_{\theta}$  from the base detector to obtain features for each of the support set images, as it can be pretrained with the base detector in normal batch training. ROIAlignV2 [23] is then used to pull the features corresponding to the location of each instance of each class. We choose to crop at the feature level rather than at the image level, as features have a larger receptive field, potentially allowing for increased global context. ROIAlignV2 produces a fixed size feature  $z_{c,i} \in \mathbb{R}^{d_f \times d_h \times d_w}$  for each object instance, with  $d_f$  being the channel dimension of the final layer of the backbone, and typically  $d_h = d_w = 7$ .

**Code Predictor Head (CPH)** Given each support sample’s extracted features  $z_{c,i}$ , hypernetwork  $\mathcal{H}_{\psi}$  predicts weights  $w_{c,i} \in \mathbb{R}^{1 \times C \times k \times k}$  and bias  $b_{c,i} \in \mathbb{R}$ , where  $C$  is the pre-




Figure 2. Hypernetwork architecture of the code predictor head.

ceding channel dimension and  $k$  is the convolutional kernel size. The code predictor head (Figure 2) consists of a shared subnetwork consisting of  $3 \times 3$  convolutional layers interleaved with group normalization [67] and ReLU activation functions, followed by a layer for predicting a weight and bias. Global average pooling after the weight and bias predictor layers is used to reduce the predicted weights to the final dimensions. While the hypernetwork is capable of predicting weights  $w_{c,i}$  of arbitrary size, we choose a kernel size of  $k = 1$  so that the generated weights can be used as either convolutional or linear layer weights, allowing compatibility with both region-based and dense detection.

**Code Process Module (CPM)** In the CPM, we aggregate the predicted parameters for all samples of a class from CPH into a single set of weights  $w_c$  and bias  $b_c$ . We found a simple average of the weights and the bias across shots to be effective:  $w_c = \frac{1}{K} \sum_{i=0}^{K-1} (w_{c,i})$  and  $b_c = \frac{1}{K} \sum_{i=0}^{K-1} (b_{c,i})$ . However, directly using the class code  $w_c$  in this form can cause gradient exploding, especially when stacking multiple convolutional layers between the input features and the final predictor head [16]. As shown in Fig. 4, gradient clipping [43] can help, but occasionally the model still does not converge well, leading to high variance in model accuracy.

Our weights  $w_c$ , as generated from the input support set features at this stage, more closely resemble feature maps than classifier weights. To this end, we want to avoid directly passing  $w_c$  to the conditional classifier. Inspired by the success of  $L^2$ -normalized feature embeddings in parameterizing classifiers for zero-shot object detection [2, 21], we explore incorporating  $L^2$ -normalization of the weights  $\frac{w_c}{\|w_c\|}$ . We normalize along the channel axis (in contrast to batch normalization [26]) to ensure weights for different classes do not interact. With normalization, learning is simplified and training is stabilized by mapping the support set features onto a unit sphere.

To ensure compatibility of the normalized weights  $\frac{w_c}{\|w_c\|}$  with a non-cosine classifier, we follow [56] and add a learnable scalar parameter  $g$ , rescaling the normalized weights as  $w_c^* = \frac{g}{\|w_c\|} w_c$ . This allows us to avoid needing to adapt the classifier in the base detector. By replacing per-class norm with a universal  $g$ , we end up with less variance across all



class weights. We found that predicting the bias counteracts this negative effect. For the bias, we further add a prior bias  $b_p = -\log((1 - \pi)/\pi)$ ,  $\pi = 0.01$  following [35] and with a scalar  $g_b$ , resulting in a final bias of  $b_c^* = g_b * b_c + b_p$ .

### 3.2. Training and Evaluation Details

To train the base detector and the hypernetwork, Sylph framework requires two sequential training stages: pretraining the base detector and learning the hypernetwork.

**Base Object Detector Pretraining** We first pretrain the base detector  $D_\phi$  with batch stochastic gradient descent on base classes  $C^b$ , optimizing for classification and bounding box regression losses. We choose FCOS as our base detector; we refer the reader to [62] for further training details. The pretraining process produces trained parameters  $\theta$  and  $\phi$ , as well as class agnostic box regression parameters  $\beta$  and class codes for the base classes  $\gamma_b = \{w_{c_b}, b_{c_b}\} \forall c_b \in C^b$ . Thus, at the conclusion of this stage, we have a detector  $D_\phi$  capable of producing bounding boxes in an image for the base classes and potentially novel classes as well.

**Meta-training** During meta-training, we create few-shot episodes of  $N$  categories by sampling a set of  $N \times (K + 1)$  image and bounding box tuples  $(I, b)$  from  $C^b$ , a support set of  $N \times K$  samples, and a query set with  $N \times 1$  samples. The query set is used as input to the base detector. Only the focal loss [35] from the classification branch is computed at this stage. The primary goal at this stage is to train the few-shot hypernetwork  $\mathcal{H}_\psi$  so that it is able to map  $S^{N \times K}$  to a new set of synthesized class codes  $\gamma_{c_b}^* = (w_{c_b}^*, b_{c_b}^*)$  for classification. We freeze the whole base object detector except the four convolutional layers in the FCOS classification subnetwork and replace its initial classifier with our conditional classifier capable of taking the synthesized class codes to make predictions on the query image features. We found that finetuning these extra convolutional layers in the base detector results in better overall performance than not finetuning them (Section 5). In preliminary experiments we found that the more components/layers we initialize from pretraining the better our final AP for base classes.

**Meta-testing** To evaluate the model’s performance across all classes, we take  $K$  shots per-class samples from the whole set and make feed-forward passes through the hypernetwork one class at a time to synthesize class codes  $\gamma_c^* = \{w_c^*, b_c^*\} \forall c \in C^b \cup C^n$ . With the generated codes, the base detector is able to perform inference at the same inference speed and behavior as a normal detector. This setup of our model is denoted as *Sylph*.

## 4. Experiments

**Datasets and Metrics** We benchmark and ablate Sylph on two datasets: COCO [36] and LVIS [22]. For COCO, we follow the split commonly used for few-shot object detection [28, 44, 65]: the 60 categories that are disjoint from

PASCAL VOC [9] are used as base classes, while the remaining 20 classes are designated novel. We report experimental results for  $K = \{1, 5, 10, 20, 30\}$  shots on the COCO minival set. For LVIS-v1, we follow the organically long-tail distribution of the dataset as proposed in [65] to produce a base-novel split. LVIS contains 405 frequent classes appearing in more than 100 images, 461 common classes with 10-100 images, and 337 rare classes with fewer than 10 images, for a total of 1203 object categories. In our experiments, we use the 337 rare classes as novel classes and the 866 frequent and common classes as base classes.

For evaluation metrics, we report mean average precision (mAP) computed on a per-split basis; we run inference for both the base and novel classes in a single pass, but we report mAP separately as different models tend to have different performances across splits. For COCO, we denote the mAPs for base and novel categories as  $AP_b$  and  $AP_n$ , respectively. Similarly, for LVIS,  $AP_r$ ,  $AP_c$ , and  $AP_f$  is the average precision aggregated across rare, common, and frequent classes, respectively. For all experiments, we report the mean and standard deviation of the mAP across five meta-testing runs. We run experiments with several pretraining strategies: (1) *Default*: the model is pretrained on ImageNet-1k [55]; (2) *Aug*: large-scale jittering (LSJ) [15] and RandAugment [5] are also applied; and (3) *All*: in addition to the aforementioned augmentations, IG-50M pretrained backbone weights from PreDet [49] are used.

**Implementation Details** For all our experiments, we use a ResNet-50 backbone [24] with a feature pyramid network (FPN) [34]. We use SGD with momentum (0.9) and weight decay ( $1e-4$ ) for all training stages. During pretraining, we set the learning rate to  $1e-2$  with a batch size of 16; we increase the batch size to 128 when data augmentation is on. During meta-training, we set the learning rate ( $lr$ ) to  $5e-4$ . We uniformly sample 3-way 5-shot tasks from the base classes, with a single query image per class.

We pretrain for 90k steps ( $\sim 11$ hrs), with an extra 30k steps for meta-learning ( $\sim 13$ hrs). The  $lr$  is decreased tenfold at steps 60k and 80k in the pretraining, and at 20k and 26k during the meta-training. Finally, we limit the number of detections per image to 100 for COCO and 300 for LVIS. We build our framework on top of Detectron2 [68]; we plan to publicly release our code upon publication.

### 4.1. Incremental Few-shot Object Detection

As the only other method designed for iFSD, we primarily compare against ONCE on both COCO and LVIS. Focusing on the **finetuning-free iFSD** evaluation protocol [44], we demonstrate the effectiveness of Sylph with several pretraining strategies. In addition, we report results for a few training-intensive FSD methods as an upper bound of our finetune-free approach, including joint-training, which is normally used for long-tailed datasets [22], and

Table 1. Benchmarking on the eval split of LVIS-v1. We use  $K = 10$  shots to infer base class codes and all available data for the rare classes ( $\leq 10$ ). Both *ONCE\** and *Sylph* predict all classes in a single pass. The base and novel data checkmarks indicate whether the data is used to update model weights during an incremental learning step.

Pretrain	Method	Base Data	Novel Data	Continuous	Re-training	AP	AP <sub>r</sub>	AP <sub>e</sub>	AP <sub>f</sub>
Default	ONCE* [44]			✓		12.9 (±0.65)	6.3 (±0.38)	11.2 (±0.60)	17.7 (±0.97)
	<b>Sylph</b>				✓	<b>18.5 (±0.12)</b> (↑5.6)	<b>10.0 (±0.17)</b>	<b>16.5 (±0.25)</b>	<b>24.3 (±0.12)</b>
	TFA-ours		✓		✓	21.0	<b>11.9</b>	17.7	28.6
	TFA* [65]	✓ ( $K$ shots)	✓		✓	<b>21.1</b>	9.1	<b>21.6</b>	25.9
	Joint-train [22]	✓	✓		✓	20.7	8.7	18.2	<b>28.8</b>
Aug	ONCE* [44]			✓		8.5 (±0.24)	6.1 (±0.31)	7.8 (±0.42)	10.2 (±0.14)
	<b>Sylph</b>			✓		<b>20.7 (±0.10)</b> (↑12.2)	<b>13.9 (±0.21)</b>	<b>19.0 (±0.19)</b>	<b>25.5 (±0.02)</b>
	TFA-ours-aug		✓		✓	<b>25.1</b>	<b>17.8</b>	22.6	<b>30.9</b>
	TFA*-aug [65]	✓ ( $K$ shots)	✓		✓	24.4	16.1	<b>24.9</b>	27.6
	Joint-train [22]	✓	✓		✓	24.3	13.3	22.7	<b>30.9</b>
All	ONCE* [44]			✓		19.4 (±0.08)	12.3 (±0.33)	18.8 (±0.28)	23.3 (±0.12)
	<b>Sylph</b>			✓		<b>24.6 (±0.10)</b> (↑5.2)	<b>16.5 (±0.34)</b>	<b>23.7 (±0.17)</b>	<b>29.1 (±0.02)</b>
	TFA-ours-aug		✓		✓	<b>27.5</b>	19.3	25.6	33.0
	TFA*-aug [65]	✓ ( $K$ shots)	✓		✓	27.0	<b>19.7</b>	<b>27.6</b>	29.6
	Joint-train [22]	✓	✓		✓	27.2	18.0	26.4	<b>33.6</b>

Table 2. Benchmarking on COCO Dataset, evaluated on minimal set. We benchmark Sylph against ONCE for  $K = 1, 5, 10$  shots, with additional  $K = 20, 30$  shots for Sylph. We also include 10-shot TFA, which finetunes on novel data. To mimic the training protocol of ONCE, we apply early-stop (at 30k steps) to Sylph pretraining (denoted Sylph-es).

Shot	Method	AP <sub>n</sub>	AP <sub>b</sub>
1	ONCE [44]	0.7	17.9
	<b>Sylph</b>	<b>0.9 (±0.11)</b>	<b>29.8 (±1.17)</b>
	<b>Sylph (All)</b>	<b>1.1 (±0.14)</b>	<b>37.6 (±1.57)</b>
5	ONCE [44]	1.0	17.9
	<b>Sylph</b>	<b>1.4 (±0.12)</b>	<b>35.5 (±0.18)</b>
	<b>Sylph (All)</b>	<b>1.5 (±0.05)</b>	<b>42.4 (±0.13)</b>
10	ONCE [44]	1.2	17.9
	<b>Sylph</b>	<b>1.6 (±0.06)</b>	<b>35.8 (±0.05)</b>
	Sylph-es	2.3	22.4
	<b>Sylph (All)</b>	<b>1.7 (±0.05)</b>	<b>42.8 (±0.07)</b>
	<b>Sylph-LVIS</b>	<b>3.8 (±0.20)</b>	37.7
	Joint-train [22]	4.0	37.7
	TFA-ours	3.6	N/A
TFA* [65]	5.7	35.9	
20	<b>Sylph</b>	<b>1.62 (±0.06)</b>	<b>36.0 (±0.08)</b>
30	<b>Sylph</b>	<b>1.65 (±0.06)</b>	<b>36.1 (±0.08)</b>

a finetuning-based method known as TFA [65].

**Finetuning-free iFSD benchmarking** We primarily compare with ONCE [44], as the most relevant method in this setting. On COCO, we compare with the reported number from [44] in Table 2. For LVIS, we re-implement ONCE with a baseline version of our code generator which has no bias prediction, no weight norm, no scaling factor  $g$  in the CPM, and no convolutional layers in the shared portion of the CPH, which effectively leaves the basic components of the hypernetwork as close to the originally-proposed ONCE as possible. We denote this version ONCE\* in Table 1.

We demonstrate that the key design choices of Sylph allow it to significantly outperform ONCE on both datasets,

across all data splits. On the large-scale dataset LVIS, *Sylph* surpasses ONCE by 8% averaged across different pretraining strategies in a fair head-to-head (no additional data augmentation or pretraining data). For the heavy data augmentation setup, *Aug*, ONCE\* struggles to converge during training, resulting in much worse performance than Sylph. In particular, we show that our method is truly able to learn novel categories from few shots without forgetting base classes. For example, with early stopping during pretraining (Sylph-es in Table 2) and  $K = 10$  shots on COCO, we achieve an  $AP_n$  twice as good as ONCE, while still surpassing it by 4 points for the base classes.

**Joint-training and finetune-based iFSD as upper bounds** For the *Joint-train* method, we follow [22] to ensure its effectiveness on the novel split in the low-data regime. In particular, we perform repeat factor sampling with the factor set to 0.001 in order to balance the sampling frequency across different classes during training. We select TFA [65] to represent finetuning-based iFSD methods. For this, we adapt the TFA [65] methodology to our FCOS-based framework, following the finetuning protocol as closely as possible. Specifically, this involves two training steps: (1) pretraining of the base detector on base classes; (2) sampling  $K = 10$  shots across both base and novel classes while freezing all layers other than the box regressor and the classifier. We finetune the regressor and train a new classifier for all classes, with base classifier parameters initialized from pretraining. We denote the adapted TFA method as *TFA\**. Additionally, we make several modification to the standard TFA to bring it closer to our setup, adjusting to an incremental batch setup [44] where the novel classes are added in a single round. In particular, in the finetuning stage, (1) only  $C^m$  is used, (2) the box regressor is kept frozen, and (3) the classifier is not initialized with any pretrained base class parameters, as we do not finetune on the base classes. In this setup, a deployed

model can be directly extended without any backbone modification to novel categories. We label this version of TFA as *TFA-ours*, as it is made possible by our framework. As our model relies on a large-scale dataset in the meta-training stage, we create a variant, *Sylph-LVIS*, which uses the LVIS dataset excluding the COCO novel classes while keeping all parts of the base detector frozen so that it is able to preserve its pretrained COCO base class codes.

We report these benchmarking results on both datasets in Tables 1 and 2. An interesting observation here is that, for all the FSD methods we benchmarked, their precision on the novel split increases significantly with an increased number of base classes. As we go from the 60 base classes in COCO to 866 base classes in LVIS, *Sylph* achieves around a 9% gain in  $AP_r$  and  $AP_n$  in Tables 1 and 2, even surpassing the gains achieved by all the variants where training and finetuning is allowed. In terms of the overall precision across all classes, *Sylph*'s performance is not far from finetuning-intensive methods: just 3 and 4 points lower compared to the best performing upper bound method on LVIS. On COCO, for example, with class augmentation in the meta-learning stage, our *Sylph-LVIS* achieves 3.8 AP, only 0.2 AP short of the joint-training approach.

**More FSD model behavior analysis** On LVIS, we benchmarked all FSD methods across three pretraining strategies, showing that all methods benefit from the use of additional augmentations and large-scale weak supervision pretraining across both novel and base classes, including *Sylph*. This behavior is highly desirable for *Sylph*, as it benefits from any improvement to the base detector.

However, for a smaller scale dataset (e.g. COCO), the boost in novel class performance for *Sylph* is less than the gain on the base classes, as shown in Table 2 for rows corresponding to *Sylph (All)*. This is related to episodic learning, where more tasks lead to improved learning compared to more per-task data. We further validate this with *Sylph-LVIS*, which has a similar amount of training data, but with more tasks; impressively, we find that *Sylph-LVIS* achieves comparable accuracy to the joint-training approach. Still, we see a large precision gap between *Sylph-LVIS* on the COCO novel split and *Sylph* on LVIS rare classes, indicating that large-scale pretraining is essential, as it results in (1) a more accurate bounding box locator, and (2) a feature extractor that can generalize better to novel classes.

Also, surprisingly, the adapted simple approach *TFA-ours* is able to achieve better precision on LVIS than its standard counterpart *TFA* with our selected base detector, with the advantage of not revisiting the base data at all.

## 5. Ablations and Further Discussion

We run all experiments in this section with the *Default* pretraining strategy unless otherwise stated.

### How does the number of base classes impact the novel

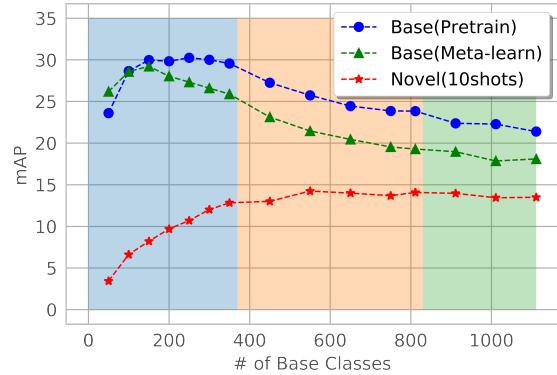


Figure 3. The effect of the number of base classes in *sylph* meta-training. The blue, orange, and green backgrounds denote frequent, common, and rare classes, respectively.

**class precision?** For this experiment, we randomly choose 50 classes from the frequent classes of LVIS as a novel set. We report base and novel mAP in Figure 3 on this fixed novel while gradually increasing the number of base classes, starting from the frequent, moving to common, and then the rare classes. For all the plotted points, we complete the training of *Sylph* on the base split, and use  $K = 10$  for class codes inference. We confirm the effect of a larger base set in novel class detection precision in Fig. 3. Indeed, novel class mAP rapidly increases in the frequent classes region, slows down when adding common classes, and starts to fully stabilize at around 800 base classes in the rare classes region. Not surprisingly, novel class score increases more rapidly when the backbone is pretrained with classes that have more samples. These results indicate for the first time that challenging incremental few-shot detection is feasible when there is a large enough base dataset.

**Model ablations** We study the effect of several key elements of our model, including the normalization scheme (GroupNorm (GN) vs  $L^2$ -Norm), the weight scaling factor  $g$ , predicted bias, and the number of convolutional layers stacked in CPH for the classifier. These results are shown in Table 3, from which we can see that overall, using either  $L^2$ -Norm or GN is very beneficial for *Sylph*, improving the overall AP across all subsets of classes by around 6 absolute points w.r.t. the baseline. However, when both GN and  $L^2$ -Norm are applied, there is no obvious extra gain. Comparing the last three rows in Table 3, we see that the use of  $g$  and bias results in a small accuracy improvement. We plotted the loss of different configurations in Fig. 4. From the leftmost figure, we can see that  $L^2$ -Norm has the largest impact in curbing the loss than any other configuration. Additionally, from the rightmost figure, we can see that both  $L^2$ -Norm and GN converge better than models without normalization. Overall, we can conclude that the elements that form *Sylph* are effective at both learning base classes and

Table 3. Ablation study: Modeling choices of *Sylph* on the LVIS.

Bias	GN	$L^2$ -Norm	$g$	#conv	AP	AP <sub>r</sub>	AP <sub>c</sub>	AP <sub>f</sub>
				0	12.9 ( $\pm 0.65$ )	6.3 ( $\pm 0.38$ )	11.2 ( $\pm 0.60$ )	17.7 ( $\pm 0.97$ )
✓				0	9.7 ( $\pm 0.25$ )	3.9 ( $\pm 0.37$ )	8.3 ( $\pm 0.27$ )	13.7 ( $\pm 0.34$ )
	✓		✓	0	17.2 ( $\pm 0.12$ )	8.8 ( $\pm 0.26$ )	15.3 ( $\pm 0.23$ )	23.2 ( $\pm 0.14$ )
✓		✓	✓	0	18.0 ( $\pm 0.12$ )	9.3 ( $\pm 0.39$ )	16.1 ( $\pm 0.15$ )	24.0 ( $\pm 0.15$ )
✓	✓		✓	0	18.0 ( $\pm 0.12$ )	9.1 ( $\pm 0.29$ )	16.2 ( $\pm 0.27$ )	24.0 ( $\pm 0.15$ )
✓	✓	✓	✓	1	18.5 ( $\pm 0.12$ )	10.0 ( $\pm 0.17$ )	16.5 ( $\pm 0.25$ )	24.3 ( $\pm 0.12$ )
✓	✓	✓	✓	2	<b>18.6 (<math>\pm 0.13</math>)</b>	<b>10.1 (<math>\pm 0.16</math>)</b>	16.7 ( $\pm 0.29$ )	<b>24.5 (<math>\pm 0.10</math>)</b>
✓	✓		✓	2	18.5 ( $\pm 0.07$ )	9.0 ( $\pm 0.17$ )	<b>16.8 (<math>\pm 0.20</math>)</b>	<b>24.5 (<math>\pm 0.08</math>)</b>
	✓	✓		2	18.4 ( $\pm 0.07$ )	9.5 ( $\pm 0.29$ )	16.7 ( $\pm 0.17$ )	24.1 ( $\pm 0.11$ )




Figure 4. The loss comparison for different model setups. On the right side, we plot the losses starting from training step 100.

generalizing to novel ones.

**Training recipe ablations** We also explored different training recipes, (1) *FA*: strictly freezing the whole base detector during meta-training, preserving the pretrained base class codes. (2) *Joint*: pretraining and meta-training on all available classes with the default setup. We report the result in Table 4. As we can see, *Sylph*’s training recipe beats that of *FA* by 2 points on the *All* setup. This means that allowing the classification convolutional subnetwork in the base detector to adapt during meta-training is important in our proposed framework. However, *Joint* performs comparably to *FA*, falling behind *Sylph* even though it has seen the novel classes during training. We think this might be explained by two reasons: (1) As the number of base classes increases, *Joint* struggles to recover the base class AP in meta-training. (2) As we use uniform sampling on the class-level, when mixed with rare classes, the more frequent classes get sampled less, thus leading to AP drop on those splits.

**Does freezing the base detector in the meta-test stage limit few-shot continual learning capabilities?** We set up a simple two step continual learning task and solve it with finetuning [42, 65]. In particular, given a pretrained FCOS model on base classes, we freeze most parts of the detector and finetune the remaining parts on all available novel data. On COCO, we keep the same base and novel splits. On LVIS, we use 100 randomly selected frequent classes as the novel split and the remaining 1103 classes as base classes. We follow *TFA\*-st* where the box regressor and

Table 4. Effect of the training recipe on the *Sylph* Framework. We report average precision across five runs.

Setup	Method	AP	AP <sub>r</sub>	AP <sub>c</sub>	AP <sub>f</sub>
Default	FA	18.3	10.5	16.3	23.9
	Joint	<b>18.5</b>	<b>10.8</b>	<b>16.5</b>	24.2
	<i>Sylph</i>	<b>18.5</b>	10.0	<b>16.5</b>	<b>24.3</b>
Aug	FA	19.9	12.5	18.5	24.8
	Joint	19.2	12.8	17.7	23.8
	<i>Sylph</i>	<b>20.7</b>	<b>13.9</b>	<b>19.0</b>	<b>25.5</b>
All	FA	22.5	15.4	21.2	27.1
	Joint	22.5	<b>17.2</b>	21.2	26.4
	<i>Sylph</i>	<b>24.6</b>	16.5	<b>23.7</b>	<b>29.1</b>

Table 5. Novel set accuracy comparison across the finetuning approaches. Full training of the detector is denoted *Scratch*.

Dataset	Setup	Method	Cls	Box	APn
COCO	All	Scratch	N/A	N/A	49.9
		TFA-ours	✓		47.1
		TFA*-st	✓	✓	50.5
LVIS	All	Scratch	N/A	N/A	34.9
		TFA-ours	✓		30.5
		TFA*-st	✓	✓	34.3

classifier are finetuned and *TFA-ours* where only the classifier is trained. The results, along with a normal training on the novel set from scratch, are shown in Table 5. We see that, surprisingly, there is no obvious performance drop for the finetune approach, and even with the strict setup *TFA-ours*, the AP<sub>n</sub> only decreases by 3 points. As *TFA-ours* closely resembles the training scheme of *Sylph*, we can conclude that the formulation we propose here does not entail a large sacrifice to the novel class learning potential.

## 6. Conclusion

We introduce *Sylph*, an object detection framework capable of extending to new classes from only a few examples in a continual manner without any training. We empirically validate that our design choices lead to effective training and improved accuracy, showing for the first time that an iFSD without test-time training can achieve performance close to finetune-based methods on large scale datasets like LVIS. While we view *Sylph* as an improvement over existing methods, there are limitations. Though we have demonstrated that pretraining a class-agnostic detector can surface novel objects with high recall, it is not infallible and still dependent on large-scale datasets. Unlabeled objects due to annotator error or a class not being in the label set can result in false negatives in the dataset, which may lead to the model failing to surface such objects [27, 30, 71]. Additionally, more sophisticated aggregation methods to fuse support set features may also lead to further improvements.

## Acknowledgements

We thank Vignesh Ramanathan, Abhijit Ogale, and Zhicheng Yan for valuable discussions and insights.



## References

- [1] Manoj Acharya, Tyler L Hayes, and Christopher Kanan. Rodeo: Replay for online object detection. *BMVC*, 2020. [2](#)
- [2] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. In *ECCV*, pages 384–400, 2018. [4](#)
- [3] Luca Bertinetto, João F Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. In *NeurIPS*, pages 523–531, 2016. [2](#)
- [4] Tung-I Chen, Yueh-Cheng Liu, Hung-Ting Su, Yu-Cheng Chang, Yu-Hsiang Lin, Jia-Fong Yeh, Wen-Chin Chen, and Winston H Hsu. Dual-awareness attention for few-shot object detection. *arXiv preprint arXiv:2102.12152*, 2021. [1](#)
- [5] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPR*, pages 702–703, 2020. [5](#)
- [6] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *NeurIPS*, pages 379–387, 2016. [2](#)
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. [1](#)
- [8] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *ICCV*, pages 6569–6578, 2019. [2](#)
- [9] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. [5](#)
- [10] Qi Fan, Wei Zhuo, Chi-Keung Tang, and Yu-Wing Tai. Few-shot object detection with attention-rpn and multi-relation detector. In *CVPR*, pages 4013–4022, 2020. [1](#)
- [11] Zhibo Fan, Yuchen Ma, Zeming Li, and Jian Sun. Generalized few-shot object detection without forgetting. In *CVPR*, pages 4527–4536, 2021. [2](#)
- [12] Sebastian Farquhar and Yarin Gal. Towards robust evaluations of continual learning. *arXiv preprint arXiv:1805.09733*, 2018. [2, 3](#)
- [13] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135. PMLR, 2017. [2](#)
- [14] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *ICML*, pages 1704–1713. PMLR, 2018. [2](#)
- [15] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, pages 2918–2928, 2021. [5](#)
- [16] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, 2018. [2, 4](#)
- [17] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. [2](#)
- [18] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *CVPR*, pages 580–587, 2014. [1, 2, 3](#)
- [19] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013. [2](#)
- [20] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, ..., Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4D: Around the World in 3,000 Hours of Egocentric Video. *CVPR*, 2022. [1](#)
- [21] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. [4](#)
- [22] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, pages 5356–5364, 2019. [1, 5, 6](#)
- [23] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017. [1, 2, 4](#)
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [1, 2, 5](#)
- [25] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, et al. Speed/accuracy trade-offs for modern convolutional object detectors. In *CVPR*, pages 7310–7311, 2017. [2, 3](#)
- [26] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456. PMLR, 2015. [4](#)
- [27] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *CVPR*, pages 5830–5840, 2021. [8](#)
- [28] Bingyi Kang, Zhuang Liu, Xin Wang, Fisher Yu, Jiashi Feng, and Trevor Darrell. Few-shot object detection via feature reweighting. In *ICCV*, pages 8420–8429, 2019. [1, 2, 4, 5](#)
- [29] Dahun Kim, Tsung-Yi Lin, Anelia Angelova, In So Kweon, and Weicheng Kuo. Learning open-world object proposals without learning to classify. *arXiv preprint arXiv:2108.06753*, 2021. [2, 4](#)
- [30] Sachin Konan, Kevin J Liang, and Li Yin. Extending one-stage detection with open-world proposals. *arXiv preprint arXiv:2201.02302*, 2022. [8](#)
- [31] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *NeurIPS*, 25:1097–1105, 2012. [1, 2](#)
- [32] Bohao Li, Boyu Yang, Chang Liu, Feng Liu, Rongrong Ji, and Qixiang Ye. Beyond max-margin: Class margin equilibrium for few-shot object detection. In *CVPR*, pages 7363–7372, 2021. [1](#)
- [33] Kevin J Liang, Samrudhdi Rangrej, Vladan Petrovic, and Tal Hassner. Few-shot learning with noisy labels. *CVPR*, 2022. [2](#)
- [34] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. [5](#)

- [35] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 2, 3, 5
- [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In *ECCV*, pages 740–755. Springer, 2014. 1, 5
- [37] Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep learning for generic object detection: A survey. *IJCV*, 128(2):261–318, 2020. 1, 2
- [38] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37. Springer, 2016. 2, 3
- [39] Michael McCloskey and Neal J Cohen. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *The Psychology of Learning and Motivation*, 1989. 1
- [40] Nikhil Mehta, Kevin J Liang, Vinay Kumar Verma, and Lawrence Carin. Continual learning using a bayesian non-parametric dictionary of weight factors. In *International Conference on Artificial Intelligence and Statistics*, pages 100–108. PMLR, 2021. 1
- [41] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 2
- [42] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. 1, 8
- [43] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML*, pages 1310–1318. PMLR, 2013. 4
- [44] Juan-Manuel Perez-Rua, Xiatian Zhu, Timothy M Hospedales, and Tao Xiang. Incremental few-shot object detection. In *CVPR*, pages 13846–13855, 2020. 1, 2, 3, 4, 5, 6
- [45] Steven Pinker. *How the mind works*. Penguin UK, 2003. 1
- [46] Hang Qi, Matthew Brown, and David G Lowe. Low-shot learning with imprinted weights. In *CVPR*, 2018. 2
- [47] Limeng Qiao, Yuxuan Zhao, Zhiyuan Li, Xi Qiu, Jianan Wu, and Chi Zhang. Defrcn: Decoupled faster r-cnn for few-shot object detection. In *ICCV*, pages 8681–8690, 2021. 1, 2
- [48] Siyuan Qiao, Chenxi Liu, Wei Shen, and Alan L Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, pages 7229–7238, 2018. 2
- [49] Vignesh Ramanathan, Rui Wang, and Dhruv Mahajan. Predet: Large-scale weakly supervised pre-training for detection. In *ICCV*, pages 2865–2875, 2021. 5
- [50] Samrudhdi Bharatkumar Rangrej, Kevin J Liang, Xi Yin, Guan Pang, Theofanis Karaletsos, Lior Wolf, and Tal Hassner. Revisiting linear decision boundaries for few-shot learning with transformer hypernetworks. 2021. 2
- [51] Roger Ratcliff. Connectionist Models of Recognition Memory: Constraints Imposed by Learning and Forgetting Functions. *Psychology Review*, 1990. 1
- [52] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. *ICLR*, 2017. 2
- [53] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016. 2
- [54] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. *NeurIPS*, 28:91–99, 2015. 2, 3
- [55] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 5
- [56] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *NeurIPS*, 29:901–909, 2016. 4
- [57] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [58] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *NeurIPS*, 30:4077–4087, 2017. 2
- [59] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *ICML*. PMLR, 2020. 1
- [60] 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*, pages 1199–1208, 2018. 2
- [61] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012. 1
- [62] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 2, 3, 4, 5
- [63] Joaquin Vanschoren. Meta-learning: A survey. *arXiv preprint arXiv:1810.03548*, 2018. 2
- [64] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *NeurIPS*, 29:3630–3638, 2016. 2
- [65] Xin Wang, Thomas E Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. *arXiv preprint arXiv:2003.06957*, 2020. 1, 2, 4, 5, 6, 8
- [66] Xin Wang, Fisher Yu, Ruth Wang, Trevor Darrell, and Joseph E Gonzalez. Tafe-net: Task-aware feature embeddings for low shot learning. In *CVPR*, pages 1831–1840, 2019. 2
- [67] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, pages 3–19, 2018. 4
- [68] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 5
- [69] Yang Xiao and Renaud Marlet. Few-shot object detection and viewpoint estimation for objects in the wild. In *ECCV*, pages 192–210. Springer, 2020. 1

- [70] Xiaopeng Yan, Ziliang Chen, Anni Xu, Xiaoxi Wang, Xiaodan Liang, and Liang Lin. Meta r-cnn: Towards general solver for instance-level low-shot learning. In *ICCV*, 2019. [1](#)
- [71] Yuewei Yang, Kevin J Liang, and Lawrence Carin. Object detection as a positive-unlabeled problem. In *BMVC*, 2020. [8](#)
- [72] Gongjie Zhang, Zhipeng Luo, Kaiwen Cui, and Shijian Lu. Meta-detr: Image-level few-shot object detection with inter-class correlation exploitation. *arXiv preprint arXiv:2103.11731*, 2021. [1](#), [2](#)