

MultIOD: Rehearsal-free Multihead Incremental Object Detector

Eden Belouadah Arnaud Dapogny Kevin Bailly

Datakalab, 143 avenue Charles de Gaulles, 92200 Neuilly-sur-Seine, France

{eb, ad, kb}@datakalab.com

Abstract

Class-Incremental learning (CIL) refers to the ability of artificial agents to integrate new classes as they appear in a stream. It is particularly interesting in evolving environments where agents have limited access to memory and computational resources. The main challenge of incremental learning is catastrophic forgetting, the inability of neural networks to retain past knowledge when learning a new one. Unfortunately, most existing class-incremental methods for object detection are applied to two-stage algorithms such as Faster-RCNN, and rely on rehearsal memory to retain past knowledge. We argue that those are not suitable in resource-limited environments, and more effort should be dedicated to anchor-free and rehearsal-free object detection. In this paper, we propose MultIOD, a class-incremental object detector based on CenterNet. Our contributions are: (1) we propose a multihead feature pyramid and multihead detection architecture to efficiently separate class representations, (2) we employ transfer learning between classes learned initially and those learned incrementally to tackle catastrophic forgetting, and (3) we use a class-wise non-max-suppression as a post-processing technique to remove redundant boxes. Results show that our method outperforms state-of-the-art methods on two Pascal VOC datasets, while only saving the model in its current state, contrary to other distillation-based counterparts.

1. Introduction

Catastrophic forgetting [30] is a significant challenge when artificial agents update their knowledge with new data. It involves losing past knowledge while rapidly transforming the model representations to fit the new data distribution. In scenarios where data arrives in streams, requiring ongoing model adaptation, the concept of Continual Learning (CL); i.e., the ability to learn from new examples while not forgetting knowledge from the old ones, gains prominence. Class-incremental learning is a subdomain of Continual Learning, where new classes are added at each system update (called state). It gained an increasing interest in the last few years

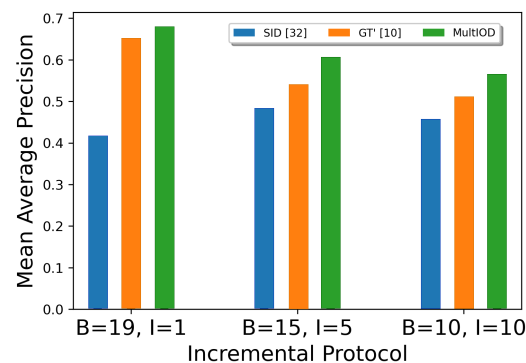


Figure 1. Mean Average Precision (IoU=0.5) on VOC0712 using different number of base classes (B) and incremental classes (I).

due to the emergence of different deep learning algorithms [1, 19, 26, 36, 40, 45]. Class-Incremental Object Detection (CIOD) is interesting in practice. It can be deployed in autonomous cars that continuously explore objects on road [39], or in security cameras to easily detect infractions, or even in large events to continuously capture attendees density for statistical purposes [7]. In computer vision, class-incremental learning is usually applied to classification, object detection, and segmentation. Rehearsal-free continual object detection and segmentation comprise an additional challenge compared to classification. The absence of annotations for objects belonging to earlier classes leads to their classification as background [31]. This phenomenon is called *background interference* (or shift) and it aggravates the effect of catastrophic forgetting.

CIOD community proposed some methods to tackle the problem [31]. However, most of them are not suitable for scenarios where real-time adaptation is required. Indeed, on the one hand, most existing CIOD models [9, 26, 33, 35, 40] are based on Faster-RCNN [38], a two-stage detection algorithm. Although this method has been proven useful for object detection, it comes at the expense of running speed. On the other hand, the largest room was booked for rehearsal-based methods, where past data is rehearsed to refresh the model representation and tackle forgetting. However, this scenario is inapplicable in cases where access to past data is impossible due to privacy issues or hardware limitations.

In this paper, we push the effort towards developing continual object detectors that are anchor-free and rehearsal-free. This is indeed the most challenging scenario, but the most useful in practice. Therefore, we propose *Multihead Incremental Object Detector (MultIOD)*, a new CIOD model based on CenterNet algorithm [51]. These are our main contributions:

1. **Architecture:** we propose a multihead feature pyramid [25] to share upsampling layers of a group of classes that appear in the same state, while a multihead detector is used to efficiently separate class representations (Subsec 3.3).

2. **Training:** we apply a transfer learning between classes learned in the initial state and classes learned incrementally to tackle catastrophic forgetting (Subsec 3.4).

3. **Inference:** we use a class-wise non-max-suppression as a post-processing technique to remove redundant boxes within the same class (Subsec 3.5).

Figure 1 shows that our method outperforms SID [34] and GT' [12], two distillation-based methods that are built on top of CenterNet. Note that our method only requires saving the model at the current state, while other state-of-the-art distillation-based counterparts need both the current and previous models to learn. Also, our method trains faster since the number of trainable parameters at each state is reduced thanks to the transfer learning scheme.

2. Related Work

Continual Learning is a hot topic of research, but most existing methods tackle classification task, and very few effort is dedicated to object detection and semantic segmentation. In this paper, we push the effort towards providing more solutions for continual object detection. We categorize existing methods into three groups as in [19].

2.1. Fine-Tuning-based approaches

Here, model parameters are continuously updated at each incremental state. ILOD [40] is one of the first works that tackled CIOD problem. Authors propose a loss that balances the interplay between predictions of new classes, while they use a distillation loss to minimize the discrepancy between responses for past classes from the previous and the current models. [8] and [14] both modify the Region Proposal Network (RPN) of a Faster-RCNN [38] to accommodate new classes. The former uses a fully connected layer along with distillation to classify proposals. The latter adds a knowledge distillation loss from a teacher model using a domain-specific dataset. Similarly to [8], the authors of [50] use distillation with a sampling strategy that helps better select proposals of foreground classes. [35] distills knowledge for proposals that have a strong relation with ground-truth boxes. [43] also uses distillation of output logits, but with a YoloV3 backbone. Other works distill intermediate features instead of logits. [9] proposes a Hint Loss

to tackle catastrophic forgetting, and uses an adaptive distillation approach that uses both RPN output and features.

Rehearsal of past class exemplars is widely used to tackle catastrophic forgetting. It was initially proposed in [36] to classify images incrementally. Its effectiveness was proven in classification on multiple datasets [5, 19]. In object detection, it was used by [13] and [39] who combined it with rehearsal and distillation, respectively. Meta-ILOD [17] uses both distillation and rehearsal to tackle forgetting. In addition, it uses a gradient-conditioning loss for the region of interest component as a regularization strategy. In [28], authors propose an attention mechanism for feature distillation and an adaptive exemplar selection strategy for rehearsal. [48] uses an external dataset and distills knowledge from two networks that learn past and new classes, respectively, to a new separate network. Alternatively, RILOD [21] helps in collecting and annotating data from search engines, and using it in incremental detection. Finally, IncDet [26] builds on Elastic Weights Consolidation (EWC) and uses a pseudo-labeling technique in addition to a Huber loss for regularization. The pseudo labeling is widely used by the community to tackle background interference. We provide in Appendix 8 a figure that explains it.

There are few CenterNet-based CIOD models [12, 34]. GT' [12] combines the ground-truth bounding boxes with predictions of the model in the previous state. The latter are first converted to bounding boxes. Then, redundant boxes are cleaned before to be distilled to the current model. Selective and Inter-related Distillation (SID) [34] is another continual detector where the authors propose a new distillation method. It consists in using Euclidean distance to measure the inter-relations between instances. This helps in considering the relationship among different instances, and thus improves the transfer. We compare *MultIOD* to both methods. Note that SID and GT' save both previous and current models in order to perform distillation, while *MultIOD* needs the current model only.

2.2. Parameter-Isolation-based approaches

This type of approaches uses a subset of the model parameters to learn a different set of classes each time.

MMN [22] freezes the most important weights in the neural network and uses the least useful ones to learn new classes. Weights importance is computed based on their magnitudes that are evaluated after each incremental state. Alternatively, [49] uses pruning techniques [29] to remove useless channels and residual blocks, while it uses a YoloV3-based ensemble network to perform detection.

MultIOD incorporates features of a parameter-isolation technique. We assign a dedicated detection head to each class while sharing a common feature pyramid among classes within the same state. The backbone remains the only component shared across the detector.

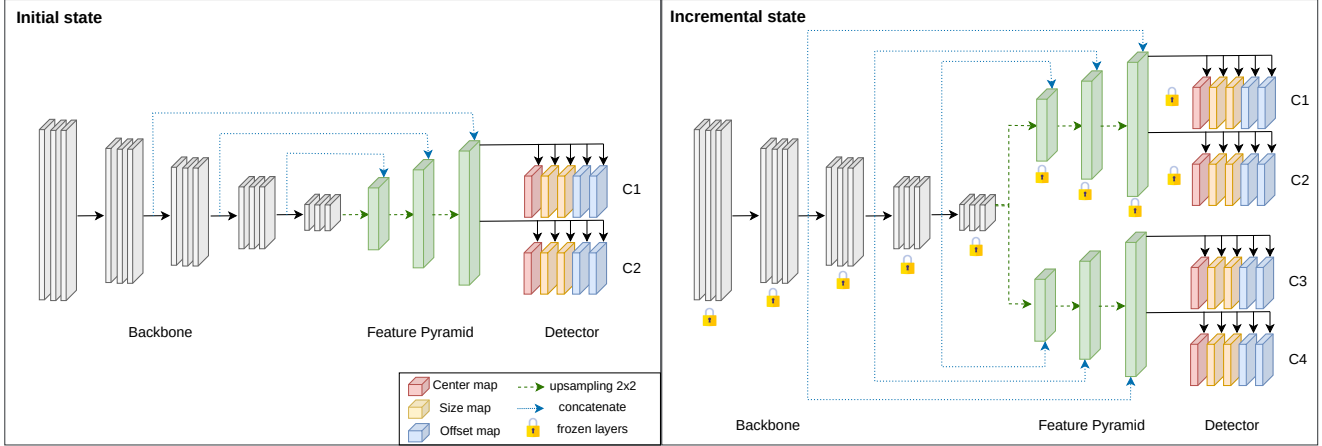


Figure 2. Illustration of *MultIOD*, depicting two toy states: one initial state (on the left), and one incremental state (on the right). The model is trained classically in the initial state using data from classes C1 and C2, while in the incremental state, it is updated using data from classes C3 and C4 only. The backbone, the feature pyramid of classes C1 and C2, as well as their detection heads are frozen once these classes are learned. Only the feature pyramid of classes C3 and C4, and their detection heads are trained in the incremental state.

2.3. Fixed-Representation-based approaches

These approaches are based on a frozen feature extractor to transfer knowledge between the initial and incremental classes. In *RODEO* [1], authors freeze the neural network after learning the first batch of classes. Later, they quantize the extracted features in order to have compact image representations. The latter are rehearsed in incremental states.

Transfer learning was proven to cope well with class-incremental learning when no memory of the past is allowed. [4] trains a deep feature extractor on initial classes and a set of Support Vector Machines (SVMs) [6] is used to learn classes incrementally. *MultIOD* shares the same spirit as [4], where the CenterNet backbone is shared between all classes, and a different detection head is specialized to learn each class. *MultIOD* is also considered as a fixed-representation-based algorithm.

Most of the works from the literature are built on top of a Faster-RCNN object detector. However, this detector cannot run in real-time. Real-life applications require a faster model to learn continuously while acquiring data in streams. According to a recent survey [3], keypoint-based detectors are the fastest, while being efficient. That is why we choose CenterNet [51] as a main algorithm to develop our class-incremental detector *MultIOD*. While multihead CenterNet has been investigated previously in [16], its primary focus was on accommodating multi-task objectives such as object detection, pose estimation, and semantic segmentation, instead of continual learning.

3. Proposed Method

3.1. Problem Definition

Let's consider a set of states $\mathcal{S} = \{S_0, S_1, \dots, S_{n-1}\}$, where n is the number of states. The initial state S_0 contains B

classes, and incremental states $S_{i>0}$ contain I classes each. In a general form, we note $|\mathcal{C}_i|$ the total number of classes seen so far in a state S_i . $\mathcal{D} = \{D_0, D_1, \dots, D_{n-1}\}$ are the sets of images of each state (these sets can overlap). An initial model M_0 is trained classically on D_0 which contains the first B classes. Incremental models M_1, M_2, \dots, M_{n-1} are trained on D_1, D_2, \dots, D_{n-1} in states S_1, S_2, \dots, S_{n-1} , respectively. Note that objects from S_i can appear in $D_{j>i}$ and objects from $D_{j>i}$ can be present in S_i but, each time, only objects from classes of the corresponding state S_i are annotated (objects from other states are not). This is known as background interference (explained in Appendix 8), a phenomenon which augments the complexity of class-incremental learning for object detection.

At each incremental state $S_{i>0}$, the model M_i is initialized with weights of model M_{i-1} . M_i has access to all training ground-truth bounding boxes from new classes, but no bounding box from past classes is available. When testing, annotations from all classes $|\mathcal{C}_i|$ are available to evaluate the performance of the model on all data seen so far. Having no access to past data during training is the most challenging scenario in class-incremental learning, yet the most interesting one in practice.

An object detection model usually comprises:

- **A Backbone:** a classification network without the final dense layer. Example: ResNet [15], EfficientNet [41], etc.
- **An Upsampling Network:** receives the output from the backbone and increases its dimensions to generate final prediction maps with enhanced resolution. Example: Deformable Convolutions [10], Feature Pyramids [25], etc.
- **A Detection Network:** takes the output of the upsampling network and makes the final prediction. Example: CenterNet [51], Yolo [37], SSD [27], etc.

Since *MultIOD* is based on CenterNet, we briefly re-

mind its main components in the next subsection.

3.2. CenterNet Object Detector

CenterNet [51] is an anchor-free object detector. It considers objects as points, and it outputs three types of maps:

- **Center map:** encodes the center of the objects using a Gaussian distribution, where the mean corresponds to the object center location, the peak value is set to 1, and the standard deviation varies based on the size of the object. The focal loss is used in [51] as an objective (Equation 1):

$$\mathcal{L}_{focal} = \frac{-1}{N} \sum_{xyc} \begin{cases} (1 - \hat{Y}_{xyc})^\alpha \log(\hat{Y}_{xyc}) & \text{if } Y_{xyc} = 1 \\ (1 - Y_{xyc})^\beta (\hat{Y}_{xyc})^\alpha & \text{otherwise} \\ \log(1 - \hat{Y}_{xyc}) & \end{cases} \quad (1)$$

where x and y are center map coordinates, c is the class index, α, β are parameters of the focal loss. N is the number of objects in the image. \hat{Y} is the predicted center map, and Y is the ground-truth center map.

- **Size map:** consists of two maps (for width and height). At the location corresponding to the peak of each Gaussian (center of the object), the width of this object is inserted in the width map, and its height in the height map. The size loss in [51] is an L1 loss (Equation 2):

$$\mathcal{L}_{size} = \frac{1}{N} \sum_{i=0}^{N-1} |\hat{S}_i - S_i| \quad (2)$$

where \hat{S} is the predicted size map, and S is the ground-truth size map. Note that the loss is only computed on map coordinates corresponding to the center of the objects.

- **Offset map:** consists also of two maps (for x and y axes). It is used to recover the discretization error caused by output stride when encoding object centers. [51] use an L1 loss, and they compute it on centers only (Equation 3):

$$\mathcal{L}_{offset} = \frac{1}{N} \sum_{i=0}^{N-1} |\hat{O}_i - O_i| \quad (3)$$

where \hat{O} is the predicted offset map, and O is the ground-truth offset map. The overall loss is the combination of the three losses (Equation 4):

$$\mathcal{L} = \lambda_{focal} \times \mathcal{L}_{focal} + \lambda_{size} \times \mathcal{L}_{size} + \lambda_{offset} \times \mathcal{L}_{offset} \quad (4)$$

where $\lambda_{focal} = 1.0$, $\lambda_{size} = 0.1$, $\lambda_{offset} = 1.0$ [51].

Note that CenterNet uses one center map per class, and shared size and offset maps between all classes.

3.3. Multihead CenterNet for Class Separability

We build on CenterNet to develop our class-incremental detector. The central concept driving *MultiIOD* is to differentiate the parameters responsible for learning distinct classes. This differentiation is not about complete separation, but rather about determining which parameters should be shared and which should be selectively distinguished. Therefore, we propose the creation of a multihead architecture, which we will elaborate on hereafter:

- **Multihead Feature Pyramid:** Feature Pyramid Network (FPN) [25] helps to build high-level semantic feature maps at all scales. This network is efficient for detecting both large and small objects. In an incremental scenario, we propose to use one feature pyramid for each group of classes that occur at the same time in the data stream. This helps in sharing a subset of parameters between these classes and reinforces their learning. The detailed architecture of our FPN is in Appendix 10. All implementation details are in Appendix 9.

- **Multihead Detector:** We adapt the original CenterNet [51] architecture to a class-incremental scenario as follows: (1) In contrast to [51] that shares size and offset maps between classes, we use one size map and one offset map for each class. The motivation behind this choice is to enable the separation of class representations in the architecture. Furthermore, the original CenterNet, by definition [51], lacks the capability to address scenarios in which two objects of distinct classes happen to share the exact same center position. This limitation results in the necessity to encode either one object or the other in the size map, but not both simultaneously.

3.4. Transfer Learning to Tackle Forgetting

In a rehearsal-free protocol, updating our model using only new data will result in a significant bias towards these specific classes in the weights of the backbone, feature pyramids, and detection heads. This bias leads to catastrophic forgetting, even when a multihead architecture is employed. To mitigate this, we propose a strategy that consists in freezing the backbone, feature pyramids of previously learned classes, and their corresponding detection heads once they are learned. This strategy efficiently minimizes the distribution shift of model parameters, facilitating effective transfer between the classes learned during the initial state and those learned incrementally. Transfer learning was proven to be effective on rehearsal-free continual learning for both classification [4] and detection [1]. It is noteworthy to mention that freezing a large part of the neural network leads to faster training since we have fewer parameters to optimize.

Figure 2 provides an overview of our method. It depicts two states: one initial, and one incremental. In the initial state, since we have only one group of classes, we use one

feature pyramid shared between these classes, and one detection head per class. We perform a classical training with all data from these classes. In the incremental state, we freeze the backbone, the feature pyramid of past classes and their detection heads. Then, we add a new feature pyramid and detection heads to learn the new classes.

Training Objective Since we use one separate size map and offset map for each class, we modify their respective losses, as in Equations 5 and 6.

$$\mathcal{L}'_{size} = \frac{1}{N} \sum_{c=0}^{|\mathcal{C}|-1} \sum_{i=0}^{N_c-1} |\hat{S}_{ic} - S_{ic}| \quad (5)$$

where $|\mathcal{C}|$ is the number of classes seen so far, and $N = N_0 + N_1 + \dots + N_{|\mathcal{C}|-1}$ is the total number of objects in the image. Similarly for the offset loss:

$$\mathcal{L}'_{offset} = \frac{1}{N} \sum_{c=0}^{|\mathcal{C}|-1} \sum_{i=0}^{N_c-1} |\hat{O}_{ic} - O_{ic}| \quad (6)$$

3.5. Class-wise Non-Max-Suppression

As mentioned in Section 3.1, the model is evaluated on all classes seen so far. In object detection, the bounding boxes generation is usually followed by a Non-Max-Suppression (NMS) operation, in which redundant boxes are eliminated, and only the most pertinent ones are kept. However, in CenterNet [51], authors demonstrate that NMS is not useful in their algorithm. The latter directly predicts object centers, resulting in a situation where each object corresponds to just one positive anchor in the ground truth.

In *MultIOD*, however, we found it very useful to use NMS and we will provide empirical evidence in Table 3 to support our finding. We select to use class-wise NMS in our model for two reasons. First, since the backbone is frozen after learning the first set of classes, the model tends to favor predictions of past classes whose confidence scores are higher than those of new classes. Second, it is important to not remove boxes belonging to different classes but share the same location. Algorithm 1 details class-wise NMS.

4. Experiments

4.1. Compared Methods

We compare *MultIOD* with CenterNet-based methods:

- **Vanilla Fine-tuning (FT)** - the lower-bound method in which we directly fine-tune CenterNet without any specific strategy to tackle catastrophic forgetting.

- **Learning without Forgetting (LwF)** [23] - the previous model is distilled to guide the training of the current model. Note that this method was initially proposed for classification.

Algorithm 1 Class-wise Non-Maximum Suppression

```

1: function CLASSWISENMS(detections, threshold)
2:   Sort detections by decreasing confidence scores
3:   within each class
4:   Initialize an empty list selected-dets
5:   for class in classes do
6:     class-dets  $\leftarrow$  detections of class
7:     while class-dets is not empty do
8:       max-det  $\leftarrow$  detection with highest
9:         confidence in class-dets
10:      Add max-det to selected-dets
11:      Remove max-det from class-dets
12:      for det in class-dets do
13:        if IoU(det, max-det)  $\geq$  threshold then
14:          Remove det from class-dets
15:   return selected-dets

```

- **SID** [34] - this is a distillation-based method, where authors distill intermediate features between the past and current model, and also distances between samples. The method was applied on both CenterNet and FCOS [42] detectors, but only the former is used here for comparability.

- **SDR** [32] - this method was initially proposed for semantic segmentation. It is based on three main components: prototype matching, feature sparsification, and contrastive learning. We try two versions of this method: SDR1 distills only the center map, while SDR2 distills all maps.

- **GT'** [12] - instead of directly distilling CenterNet maps from the previous to the current model, the past model is used to annotate data from new classes. After extracting the bounding boxes, the latter are encoded along with the ground truth before training the model on both the ground-truth and the pseudo-labeled data. Please note that the experimental protocol of this method is not realistic because authors remove images where past class objects are present.

- **Full** - the upper-bound method. A full training is performed on all classes with all their available data.

Note that we take results of *LwF* and *SDR* from [34] and [32], respectively.

For a fair comparison, we only compare our method with algorithms that do not perform any rehearsal, but we still report results against some two-stage continual detectors that use or not rehearsal, in Appendix 13. While all the aforementioned methods build upon CenterNet with a ResNet50 backbone, we opt to use the EfficientNet family of models as backbone because they are more suitable for efficient training. Unless specified otherwise, our choice of backbone is EfficientNet-B3 because it provides the best trade-off between performance and number of parameters. CenterNet with ResNet50 contains about 32M parameters, while CenterNet with EfficientNet-B3 model contains 17.7M parameters only. That is, we reduce around

45% of the number of parameters. We will prove in the results section, that even though our model achieves lower results in classical training (with all data available at once), it still outperforms by a good margin the other methods in almost all configurations. Our model is not only more robust against forgetting, but it is also parameter-efficient.

4.2. Datasets

- **MNIST-Detection** (MNIST below) - this dataset was initially designed for digit classification [20]. We create an object-detection toy version of it to perform our ablation studies as it runs faster than other datasets. We used this github repository ¹ to create training and validation sets of 5000 and 1000 images of size 512×512, respectively. We made sure to create a challenging dataset. The dataset creation procedure is detailed in Appendix 11.

- **Pascal VOC2007** [11] - this dataset contains 20 object classes with 5,011 and 4,952 training and validation examples, respectively. Note that we use both the training and validation sets for training, as in [34].

- **Pascal VOC0712** [11] - here we use both the training and validation sets of VOC 2007 and 2012 as in [12, 34, 51]. We use the test set of VOC2007 for validation. In total, this dataset contains 16,550 and 4,952 training and validation images, respectively.

4.3. Methodology

We first apply most of the ablations on MNIST dataset, then perform the final experiments on VOC2007 and VOC0712. Following [17, 31, 34], We evaluate our method using three incremental learning scenarios: we order classes (numerically for MNIST, alphabetically for VOC), then divide them into two states. For MNIST, we use $B = 9, I = 1$; $B = 7, I = 3$ and $B = 5, I = 5$. We use the following number of classes per state for the VOC dataset: $B = 19, I = 1$; $B = 15, I = 5$ and $B = 10, I = 10$. Varying the number of classes between states is important to assess the robustness of CL methods.

4.4. Evaluation Metrics

- **mAP@0.5** - the mean average precision computed at an IoU (Intersection-over-Union) threshold of 0.5.

- **mAP@[0.5, 0.95]** - the mean average precision averaged on IoU threshold that varies between 0.5 and 0.95 with a step of 0.05. Note that results with this metric are presented in Appendix 12.

- **F_{mAP}** - the harmonic mean between the mean average precision of past and new classes.

$$F_{mAP} = 2 \times \frac{mAP_{past} \times mAP_{new}}{mAP_{past} + mAP_{new}} \quad (7)$$

¹<https://github.com/hukkelas/MNIST-ObjectDetection>

When $F_{mAP} = 0.0$, this means that the model completely failed to detect either past classes or new classes.

5. Ablation Studies

5.1. Ablation of Backbones

The backbone has a direct impact on model performance because it is responsible for the quality of extracted features. In Table 1, we vary the backbone and use EfficientNet-B0 (10 million params), EfficientNet-B3 (17.7 million params), and EfficientNet-B5 (37 million params) on VOC2007 and VOC0712. We provide results on MNIST in Appendix 13. Results show that the best overall backbone is EfficientNet-B3. EfficientNet-B0 is smaller and has more difficulties to generalize. However, it still provides decent performances while having 10M params only. EfficientNet-B5 provides comparable results with EfficientNet-B3. We decided to keep the latter since it has the best compromise between parameters number and performance.

5.2. Ablation of Frozen Layers

One of the main components of our method is to use transfer learning between classes learned in the initial state, and classes that are learned incrementally. Therefore, studying the impact of freezing the different layers of our system is a crucial part to assess its robustness. In Table 2, we use MNIST dataset and EfficientNet-B0 backbone in a setting where $B = 5$ and $I = 5$, and test the following configurations: (1) do not freeze anything in the architecture, (2) freeze the backbone only, (3) freeze both the backbone and the feature pyramid of past classes, and (4) our full method, in which we freeze not only the backbone, but also the feature pyramid of past classes, and their detection heads. Results show that it is crucial to freeze past class detection heads in order to avoid their catastrophic forgetting. In fact, it is meaningful to freeze them only if all the preceding components are also frozen (the feature pyramid of past classes and the backbone).

backbone	feature pyramid	detection head	mAP	F _{mAP}
×	×	×	44.2	0.0
🔒	×	×	44.3	0.0
🔒	🔒	×	42.0	0.0
🔒	🔒	🔒	91.3	91.3

Table 2. Performance of *MultiOD* on MNIST with EfficientNet-B0 and $B = 5, I = 5$ when ablating its components freezing.

5.3. Ablation of Non-Max-Suppression Strategies

Using multiple detection heads leads to the risk of predicting multiple objects (from different classes) at the same pixel location. Thus, it is important to eliminate irrelevant bounding boxes (false positives). In this experiment, we ablate the following NMS strategies:

- **No-NMS**: this is the method originally used in CenterNet [51], where no non-max-suppression is applied.

Method	VOC2007							VOC0712						
	Full	$B = 19, I = 1$		$B = 15, I = 5$		$B = 10, I = 10$		Full	$B = 19, I = 1$		$B = 15, I = 5$		$B = 10, I = 10$	
	mAP	mAP	F_{mAP}	mAP	F_{mAP}	mAP	F_{mAP}	mAP	mAP	F_{mAP}	mAP	F_{mAP}	mAP	F_{mAP}
EfficientNet-B0	56.7	55.7	35.6	49.2	33.8	46.3	45.5	66.9	65.0	54.0	57.7	44.7	53.9	52.9
EfficientNet-B3	60.4	59.9	39.2	52.6	37.8	48.4	47.2	69.5	68.0	56.9	60.7	47.0	56.6	55.8
EfficientNet-B5	61.3	59.9	45.2	52.5	36.5	45.2	42.4	70.3	68.4	54.7	60.8	46.5	55.3	53.9

Table 1. Ablation of backbones on VOC2007 and VOC0712 datasets (mAP@0.5).

- Inter-class NMS: here, a standard NMS algorithm is utilized to eliminate redundant bounding boxes, irrespective of their class membership.

- Class-wise NMS: as described in Subsection 3.5, it is designed to enhance the precision of boxes selection within individual classes, effectively curbing redundancy and selecting the most pertinent boxes for each specific class.

- Soft-NMS: unlike traditional NMS, which remove boxes that overlap significantly, it applies a decay function to the confidence scores of neighboring boxes, gradually reducing their impact. This results in a smoother suppression of redundant boxes and helps retain boxes with slightly lower scores that might still contribute to accurate detection.

Table 3 shows that the method which provides the best results is class-wise NMS, the one we choose to use. The second best method is the standard inter-class NMS. The former helps in detecting more objects, specifically when objects of different classes are highly overlapped. This comes at the expense of having some false positives. In contrast, inter-class NMS reduces the number of false positives at the expense of not detecting bounding boxes of different classes that are at the same location in the image. Depending on the use case, one or the other method can be preferred. In our use case, Soft-NMS did not provide good results. This can be explained by the fact that the gradual decay of confidence scores in Soft-NMS might inadvertently allow boxes with lower scores to persist, potentially leading to false positives or less accurate detections. Finally, unsurprisingly, the method with the worst performance is No-NMS due to the high number of false positives. Note that the same ablation for VOC2007 is in Appendix 14.

Method	Full	$B = 9, I = 1$		$B = 7, I = 3$		$B = 5, I = 5$	
	mAP	mAP	F_{mAP}	mAP	F_{mAP}	mAP	F_{mAP}
No-NMS	90.2	89.5	89.7	90.8	91.3	88.4	88.2
Soft-NMS	91.5	89.7	89.5	91.1	91.4	88.5	88.4
Inter-class NMS	91.8	90.2	90.2	91.7	92.0	89.9	89.8
Class-wise NMS	93.1	91.3	91.2	93.1	93.5	91.3	91.3

Table 3. Performance of our model using MNIST dataset with different NMS strategies and EfficientNet-B0.

6. Results and Discussion

Tables 4 and 5 provide the main results of *MultIOD* compared to state-of-the-art methods on VOC2007 and VOC0712 datasets, respectively.

Results show that vanilla Fine-tuning behaves the worst as it pushes for the complete plasticity of the neural net-

Method	Full	$B = 19, I = 1$		$B = 15, I = 5$		$B = 10, I = 10$	
	mAP	mAP	F_{mAP}	mAP	F_{mAP}	mAP	F_{mAP}
FT		7.3	9.0	13.2	5.0	28.6	0.0
LwF [23]	65.6	12.5	10.9	6.9	1.5	17.6	0.0
SID [34]		45.5	27.0	51.9	46.1	43.3	43.0
Ours	60.4	59.9	39.2	52.6	37.8	48.4	47.2

Table 4. mAP@0.5 and F_{mAP} scores on VOC2007 dataset.

work. This is an extreme case in which the mAP for past classes is equal to zero or nearly so, which causes the F_{mAP} score to be zero as we can see in the case of $B = 10, I = 10$.

LwF [23] is a method that was initially proposed to tackle continual learning in classification task. It consists in distilling the current model from the previous one, in order to mimic its predictions. Even though this method provides good results in rehearsal-free continual classification [5], it fails to generalize on the detection task. One reason could be that classification task is orders of magnitude easier than object detection, and more specialized techniques are needed in order to tackle the latter.

Our method outperforms results from state of the art in 11 cases out of 12. For VOC2007 dataset, in the scenario $B = 19, I = 1$, our method gains up to 14.4 mAP points compared to the state-of-the-art method SID [34]. This is intuitive insofar as fixed-representations are more powerful when trained with more classes in the initial state. The diversity of the initial dataset and its size is crucial to have universal representation and transfer to unseen classes.

In the scenario $B = 15, I = 5$, *MultIOD* outperforms SID in terms of mAP score, but fails to do so for the F_{mAP} score. This indicates that in this specific case, SID balances better the compromise between plasticity (the ability to adapt to new class distribution), and stability (the ability to keep as much knowledge as possible from the past).

Method	Full	$B = 19, I = 1$		$B = 15, I = 5$		$B = 10, I = 10$	
	mAP	mAP	F_{mAP}	mAP	F_{mAP}	mAP	F_{mAP}
FT		20.2	26.3	16.3	17.3	28.0	0.0
LwF [23]		36.3	28.1	12.0	11.5	22.5	0.0
SDR 1 [32]	73.1	60.4	22.8	41.4	23.9	35.5	32.2
SDR 2 [32]		49.8	25.8	21.9	22.3	30.1	9.6
SID [34]		41.6	19.3	48.4	17.9	45.5	44.2
GT [12]		65.2	39.2	54.0	33.8	51.2	50.9
Ours	69.5	68.0	56.9	60.7	47.0	56.6	55.8

Table 5. mAP@0.5 and F_{mAP} scores on VOC0712 dataset.

For VOC0712, *MultIOD* outperforms other methods in

all cases, followed by GT' [12]. It is noteworthy to mention that the experimental protocol of this method is not realistic insofar as authors remove images that contain objects from past classes when training the model on new classes. Thus, they avoid the problem of background shift by removing the overlap between annotations of background and past class objects. However, in real-life situations, it is impossible to have this separation as both past and new class objects can appear in the scene, and the model should be able to effectively learn new classes while past ones are also there. SDR [32] is a method that was initially proposed for continual semantic segmentation, and was later adapted by [12] to object detection. Regardless of the effectiveness of this method on semantic segmentation, it does not achieve the same impressive performance on continual object detection. It reaches good scores in the protocol $B = 19, I = 1$ only. One reason for that could be that in CenterNet there is less information to distill compared to segmentation where all object pixels are used.

Coincidentally, all CenterNet-based continual object detectors to which we compare *MultIOD*, use distillation to update model weights from its past states. Overall, results show the usefulness of transfer learning compared to distillation. This is true especially when classes learned incrementally belong to the same domain as classes learned initially.

In both Tables 4 and 5, it is crucial to emphasize that our upper-bound (Full) model scores lower by 5.2 and 3.6 mAP points respectively, compared to the classical training from the state of the art. Nevertheless, it is interesting to notice that even with this drop in performance under traditional training, our model excels in comparison to other methods across nearly all scenarios involving incremental learning. This finding proves that *MultIOD* is more robust against catastrophic forgetting, because the gap with our upper-bound model is tighter compared to the other methods with their respective upper-bound model.

We remind that the difference in performance between the two upper-bound models is that our detector is built on top of EfficientNet-B3 backbone, while the other methods use a ResNet50. In addition to the reduction in the number of parameters (17.7M vs. 32M), *MultIOD* does not require keeping the past model to learn new classes. In contrast, all methods used for comparison need the past model in order to extract its activation and distill its knowledge to the current model. Thus, these methods require keeping two models simultaneously to make the training happen. Some predictions with *MultIOD* are in Appendix 15.

6.1. Additional Experiment

In this experiment, we take the challenge to a more difficult level by adding only one or two classes in each incremental state. In this case, we use one feature pyramid for each class

to have a complete separation of their representations in the upsampling and detection heads. To avoid an explosion in the number of parameters compared to the previous protocol, we reduce the number of filters in the feature pyramids to have a comparable number of parameters with the model used in previous experiments. Details of this architecture are in Appendix 10. Results in Table 6 show that we lose from 1 to 6 mAP points approximately when passing from one state to another. Unfortunately, CenterNet-based state-of-the-art methods do not use this protocol, and we thus cannot compare them with *MultIOD*. We provide results for future use.

States	S_0	S_1	S_2	S_3	S_4	S_5
	$B = 15$		$I = 1$			
FT	56.6	0.7	0.3	1.0	0.3	1.8
Ours		53.4	47.9	45.6	43.3	42.7
	$B = 10$		$I = 2$			
FT	57.7	5.3	4.7	5.1	2.1	3.2
Ours		48.4	45.7	43.0	39.7	38.0

Table 6. mAP@0.5 of *MultIOD* on VOC2007.

7. Conclusions

In this paper, we present *MultIOD*, a class-incremental object detector based on CenterNet [51]. Our approach uses a multihead detection component, along with a frozen backbone. A multihead feature pyramid is also used to ensure a satisfying trade-off between stability and plasticity. Finally, it involves an efficient class-wise NMS that provides robustness to remove duplicate bounding boxes. Results show the effectiveness of our approach against CenterNet-based methods on Pascal VOC datasets [11] in many incremental scenarios. However, *MultIOD* has some limitations:

- **Scalability** - One of the main drawbacks of our method is its limited capacity to scale, since we add one detection head for each incremental class. It would be interesting to investigate more intelligent ways to regroup classes by their semantic similarities, and create new detection heads only when the new classes are significantly different from already existing ones. This can be also combined with a gating mechanism [2] that automatically selects relevant head(s) to use during inference.

- **Quality of the fixed representation** - This drawback is common to all techniques that rely on transfer learning. If the initial network is trained on sufficiently rich data, the quality of the transfer is prominent, but inversely, if the fixed representation is poor, the learning of subsequent classes drastically drops in performance. A good perspective would be to build more universal representations [44] and test the transferability of better datasets, such as COCO [24].

Acknowledgment: This work was granted access to the HPC resources of IDRIS under the allocation 2023-A0141014124 made by GENCI.

References

- [1] Manoj Acharya, Tyler L. Hayes, and Christopher Kanan. RODEO: replay for online object detection. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020. 1, 3, 4
- [2] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 7120–7129. IEEE Computer Society, 2017. 8
- [3] Elahe Arani, Shruthi Gowda, Ratnajit Mukherjee, Omar Magdy, Senthilkumar Kathiresan, and Bahram Zonooz. A comprehensive study of real-time object detection networks across multiple domains: A survey. *CoRR*, abs/2208.10895, 2022. 3
- [4] Eden Belouadah and Adrian Popescu. Deesil: Deep-shallow incremental learning. In *Computer Vision - ECCV 2018 Workshops - Munich, Germany, September 8-14, 2018, Proceedings, Part II*, pages 151–157. Springer, 2018. 3, 4
- [5] Eden Belouadah, Adrian Popescu, and Ioannis Kanellos. A comprehensive study of class incremental learning algorithms for visual tasks. *CoRR*, abs/2011.01844, 2020. 2, 7
- [6] Bernhard E. Boser, Isabelle Guyon, and Vladimir Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual ACM Conference on Computational Learning Theory, COLT 1992, Pittsburgh, PA, USA, July 27-29, 1992*, pages 144–152. ACM, 1992. 3
- [7] Mayur D Chaudhari and Archana S Ghotkar. A study on crowd detection and density analysis for safety control. *International journal of computer sciences and engineering*, 6: 424–428, 2018. 1
- [8] Jingzhou Chen, Shihao Wang, Ling Chen, Haibin Cai, and Yuntao Qian. Incremental detection of remote sensing objects with feature pyramid and knowledge distillation. *IEEE Trans. Geosci. Remote. Sens.*, 60:1–13, 2022. 2
- [9] Li Chen, Chunyan Yu, and Lvcai Chen. A new knowledge distillation for incremental object detection. In *International Joint Conference on Neural Networks, IJCNN 2019 Budapest, Hungary, July 14-19, 2019*, pages 1–7. IEEE, 2019. 1, 2
- [10] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 764–773. IEEE Computer Society, 2017. 3
- [11] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010. 6, 8, 1
- [12] Sumyung Gang, Daewon Chung, and Joonjae Lee. Predictive distillation method of anchor-free object detection model for continual learning. *Applied Sciences*, 12(13):6419, 2022. 2, 5, 6, 7, 8, 1
- [13] Yu Hao, Yanwei Fu, and Yu-Gang Jiang. Take goods from shelves: A dataset for class-incremental object detection. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval, ICMR 2019, Ottawa, ON, Canada, June 10-13, 2019*, pages 271–278. ACM, 2019. 2
- [14] Yu Hao, Yanwei Fu, Yu-Gang Jiang, and Qi Tian. An end-to-end architecture for class-incremental object detection with knowledge distillation. In *IEEE International Conference on Multimedia and Expo, ICME 2019, Shanghai, China, July 8-12, 2019*, pages 1–6. IEEE, 2019. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 3
- [16] Falk Heuer, Sven Mantowsky, Syed Saqib Bukhari, and Georg Schneider. Multitask-centernet (MCN): efficient and diverse multitask learning using an anchor free approach. In *IEEE/CVF International Conference on Computer Vision Workshops, ICCVW 2021, Montreal, BC, Canada, October 11-17, 2021*, pages 997–1005. IEEE, 2021. 3
- [17] K. J. Joseph, Jathushan Rajasegaran, Salman H. Khan, Fahad Shahbaz Khan, Vineeth Balasubramanian, and Ling Shao. Incremental object detection via meta-learning. *CoRR*, abs/2003.08798, 2020. 2, 6
- [18] K. J. Joseph, Salman H. Khan, Fahad Shahbaz Khan, and Vineeth N. Balasubramanian. Towards open world object detection. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 5830–5840. Computer Vision Foundation / IEEE, 2021. 3
- [19] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory G. Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(7):3366–3385, 2022. 1, 2
- [20] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. <http://yann.lecun.com/exdb/mnist/>, 2010. 6, 1
- [21] Dawei Li, Serafettin Tasci, Shalini Ghosh, Jingwen Zhu, Junting Zhang, and Larry P. Heck. RILOD: near real-time incremental learning for object detection at the edge. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing, SEC 2019, Arlington, Virginia, USA, November 7-9, 2019*, pages 113–126. ACM, 2019. 2
- [22] Wei Li, Qingbo Wu, Linfeng Xu, and Chao Shang. Incremental learning of single-stage detectors with mining memory neurons. In *2018 IEEE 4th International Conference on Computer and Communications (ICCC)*, pages 1981–1985. IEEE, 2018. 2
- [23] Zhizhong Li and Derek Hoiem. Learning without forgetting. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, pages 614–629. Springer, 2016. 5, 7
- [24] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, pages 740–755. Springer, 2014. 8
- [25] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid

- networks for object detection. *CoRR*, abs/1612.03144, 2016. [2](#), [3](#), [4](#)
- [26] Liyang Liu, Zhanghui Kuang, Yimin Chen, Jing-Hao Xue, Wenming Yang, and Wayne Zhang. Incdet: In defense of elastic weight consolidation for incremental object detection. *IEEE Trans. Neural Networks Learn. Syst.*, 32(6):2306–2319, 2021. [1](#), [2](#), [3](#)
- [27] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. *CoRR*, abs/1512.02325, 2015. [3](#)
- [28] Xialei Liu, Hao Yang, Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Multi-task incremental learning for object detection. *arXiv preprint arXiv:2002.05347*, 2020. [2](#)
- [29] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. *CoRR*, abs/1708.06519, 2017. [2](#)
- [30] Michael McCloskey and Neil J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *The Psychology of Learning and Motivation*, 24: 104–169, 1989. [1](#)
- [31] Angelo G. Menezes, Gustavo de Moura, Cézanne Alves, and André C. P. L. F. de Carvalho. Continual object detection: A review of definitions, strategies, and challenges. *CoRR*, abs/2205.15445, 2022. [1](#), [6](#)
- [32] Umberto Michieli and Pietro Zanuttigh. Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 1114–1124. Computer Vision Foundation / IEEE, 2021. [5](#), [7](#), [8](#)
- [33] Can Peng, Kun Zhao, and Brian C Lovell. Faster ilod: Incremental learning for object detectors based on faster rcnn. *Pattern Recognition Letters*, 2020. [1](#), [3](#)
- [34] Can Peng, Kun Zhao, Sam Maksoud, Meng Li, and Brian C. Lovell. SID: incremental learning for anchor-free object detection via selective and inter-related distillation. *Comput. Vis. Image Underst.*, 210:103229, 2021. [2](#), [5](#), [6](#), [7](#), [1](#)
- [35] Kandan Ramakrishnan, Rameswar Panda, Quanfu Fan, John Henning, Aude Oliva, and Rogério Feris. Relationship matters: Relation guided knowledge transfer for incremental learning of object detectors. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, June 14-19, 2020*, pages 1009–1018. Computer Vision Foundation / IEEE, 2020. [1](#), [2](#)
- [36] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *Conference on Computer Vision and Pattern Recognition*, 2017. [1](#), [2](#)
- [37] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 779–788. IEEE Computer Society, 2016. [3](#)
- [38] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99, 2015. [1](#), [2](#)
- [39] Jeng-Lun Shieh, Qazi Mazhar ul Haq, Muhamad Amirul Haq, Said Karam, Peter Chondro, De-Qin Gao, and Shanq-Jang Ruan. Continual learning strategy in one-stage object detection framework based on experience replay for autonomous driving vehicle. *Sensors*, 20(23):6777, 2020. [1](#), [2](#)
- [40] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 3420–3429. IEEE Computer Society, 2017. [1](#), [2](#)
- [41] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 6105–6114. PMLR, 2019. [3](#), [2](#)
- [42] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: fully convolutional one-stage object detection. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 9626–9635. IEEE, 2019. [5](#)
- [43] Qazi Mazhar ul Haq, Shanq-Jang Ruan, Muhamad Amirul Haq, Said Karam, Jeng-Lun Shieh, Peter Chondro, and De-Qin Gao. An incremental learning of yolov3 without catastrophic forgetting for smart city applications. *IEEE Consumer Electron. Mag.*, 11(5):56–63, 2022. [2](#)
- [44] Zhenyu Wang, Yali Li, Xi Chen, Ser-Nam Lim, Antonio Torralba, Hengshuang Zhao, and Shengjin Wang. Detecting everything in the open world: Towards universal object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 11433–11443. IEEE, 2023. [8](#)
- [45] Dongbao Yang, Yu Zhou, Wei Shi, Dayan Wu, and Weiping Wang. RD-IOD: two-level residual-distillation-based triple-network for incremental object detection. *ACM Trans. Multimed. Comput. Commun. Appl.*, 18(1):18:1–18:23, 2022. [1](#), [3](#)
- [46] Dongbao Yang, Yu Zhou, Aoting Zhang, Xurui Sun, Dayan Wu, Weiping Wang, and Qixiang Ye. Multi-view correlation distillation for incremental object detection. *Pattern Recognit.*, 131:108863, 2022. [3](#)
- [47] Shuo Yang, Peize Sun, Yi Jiang, Xiaobo Xia, Ruiheng Zhang, Zehuan Yuan, Changhu Wang, Ping Luo, and Min Xu. Objects in semantic topology. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [3](#)
- [48] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry P. Heck, Heming Zhang, and C.-C. Jay Kuo. Class-incremental learning via deep model consolidation. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, March 1-5, 2020*, pages 1120–1129. IEEE, 2020. [2](#)

- [49] Nan Zhang, Zhigang Sun, Kai Zhang, and Li Xiao. Incremental learning of object detection with output merging of compact expert detectors. In *2021 4th International Conference on Intelligent Autonomous Systems (ICoIAS)*, pages 1–7. IEEE, 2021. [2](#)
- [50] Wang Zhou, Shiyu Chang, Norma E. Sosa, Hendrik F. Hamann, and David D. Cox. Lifelong object detection. *CoRR*, abs/2009.01129, 2020. [2](#)
- [51] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *CoRR*, abs/1904.07850, 2019. [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [1](#)