

Continual Detection Transformer for Incremental Object Detection

Yaoyao Liu¹ Bernt Schiele¹ Andrea Vedaldi² Christian Rupprecht²

¹Max Planck Institute for Informatics, Saarland Informatics Campus

²Visual Geometry Group, Department of Engineering Science, University of Oxford

{yaoyao.liu, schiele}@mpi-inf.mpg.de {vedaldi, chrisr}@robots.ox.ac.uk

Abstract

Incremental object detection (IOD) aims to train an object detector in phases, each with annotations for new object categories. As other incremental settings, IOD is subject to catastrophic forgetting, which is often addressed by techniques such as knowledge distillation (KD) and exemplar replay (ER). However, KD and ER do not work well if applied directly to state-of-the-art transformer-based object detectors such as Deformable DETR [59] and UP-DETR [9]. In this paper, we solve these issues by proposing a Continual DEtection TRansformer (CL-DETR), a new method for transformer-based IOD which enables effective usage of KD and ER in this context. First, we introduce a Detector Knowledge Distillation (DKD) loss, focusing on the most informative and reliable predictions from old versions of the model, ignoring redundant background predictions, and ensuring compatibility with the available ground-truth labels. We also improve ER by proposing a calibration strategy to preserve the label distribution of the training set, therefore better matching training and testing statistics. We conduct extensive experiments on COCO 2017 and demonstrate that CL-DETR achieves state-of-the-art results in the IOD setting.¹

1. Introduction

Humans inherently learn in an incremental manner, acquiring new concepts over time without forgetting previous ones. In contrast, machine learning suffers from *catastrophic forgetting* [21, 35, 36], where learning from non-i.i.d. data can override knowledge acquired previously. Unsurprisingly, forgetting also affects object detection [2, 12, 20, 37, 44, 50, 54]. In this context, the problem was formalized by Shmelkov *et al.* [44], who defined an incremental object detection (IOD) protocol, where the training samples for different object categories are observed in phases, re-

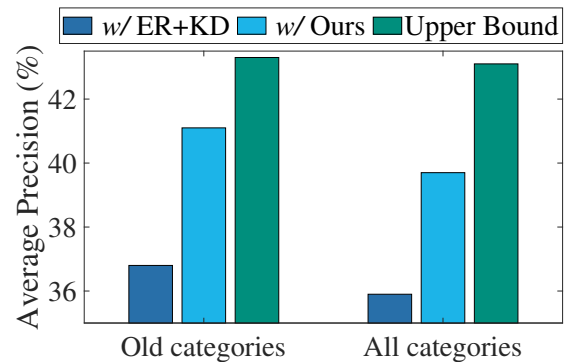


Figure 1. The final Average Precision (AP, %) of two-phase incremental object detection on COCO 2017. We observe 70 and 10 categories in the first and second phases, respectively. The baseline is Deformable DETR [59]. “Upper bound” shows the results of joint training with all previous data accessible in each phase.

stricting the ability of the trainer to access past data.

Popular methods to address forgetting in tasks other than detection include Knowledge Distillation (KD) and Exemplar Replay (ER). KD [11, 16, 17, 26, 57] uses regularization in an attempt to preserve previous knowledge when training the model on new data. The key idea is to encourage the new model’s logits or feature maps to be close to those of the old model. ER methods [5, 29, 32, 33, 41, 52] work instead by memorising some of the past training data (the *exemplars*), replaying them in the following phases to “remember” the old object categories.

Recent state-of-the-art results in object detection have been achieved by a family of transformer-based architectures that include DETR [4], Deformable DETR [59] and UP-DETR [9]. In this paper, we show that KD and ER do not work well if applied directly to these models. For instance, in Fig. 1 we show that applying KD and ER to Deformable DETR leads to much worse results compared to training with all data accessible in each phase (*i.e.*, the standard non-incremental setting).

We identify two main issues that cause this drop in per-

¹Code: <https://lyy.mpi-inf.mpg.de/CL-DETR/>

formance. First, transformer-based detectors work by testing a large number of object hypotheses in parallel. Because the number of hypotheses is much larger than the typical number of objects in an image, most of them are negative, resulting in an unbalanced KD loss. Furthermore, because both old and new object categories can co-exist in any given training image, the KD loss and regular training objective can provide contradictory evidence. Second, ER methods for image classification try to sample the same number of exemplars for each category. In IOD, this is not a good strategy because the true object category distribution is typically highly skewed. Balanced sampling causes a mismatch between the training and testing data statistics.

In this paper, we solve these issues by proposing *Continual DETection TRansformer* (CL-DETR), a new method for transformer-based IOD which enables effective usage of KD and ER in this context. CL-DETR introduces the concept of *Detector Knowledge Distillation* (DKD), selecting the most confident object predictions from the old model, merging them with the ground-truth labels for the new categories while resolving conflicts, and applying standard joint bipartite matching between the merged labels and the current model’s predictions for training. This approach subsumes the KD loss, applying it only for foreground predictions correctly matched to the appropriate model’s hypotheses. CL-DETR also improves ER by introducing a new calibration strategy to preserve the distribution of object categories observed in the training data. This is obtained by carefully engineering the set of exemplars remembered to match the desired distribution. Furthermore, each phase consists of a main training step followed by a smaller one focusing on better calibrating the model.

We also propose a more realistic variant of the IOD benchmark protocol. In previous works [12, 44], in each phase, the incremental detector is allowed to observe all images that contain a certain type of object. Because images often contain a mix of object classes, both old and new, this means that the same images can be observed in different training phases. This is incompatible with the standard definition of incremental learning [16, 33, 41] where, with the exception of the examples deliberately stored in the exemplar memory, the images observed in different phases do not repeat. We redefine the IOD protocol to avoid this issue.

We demonstrate CL-DETR by applying it to different transformer-based detectors including Deformable DETR [59] and UP-DETR [9]. As shown in Fig. 1, our results on COCO 2017 show that CL-DETR leads to significant improvements compared to the baseline, boosting AP by 4.2 percentage points compared to a direct application of KD and ER to the underlying detector model. We further study and justify our modelling choices via ablations.

To summarise, we make **four contributions**: (1) The DKD loss that improves KD for knowledge distillation by

resolving conflicts between distilled knowledge and new evidence and by ignoring redundant background detections; (2) A calibration strategy for ER to match the stored exemplars to the training set distribution; (3) A revised IOD benchmark protocol that avoids observing the same images in different training phases; (4) Extensive experiments on COCO 2017, including state-of-the-art results, an in-depth ablation study, and further visualizations.

2. Related Work

Incremental learning. Incremental learning (also known as continual learning [2, 10, 34] and lifelong learning [1, 6, 8]) aims at learning models in phases that focus on different subsets of the label space. Recent incremental learning methods can be divided into two categories: (i) Knowledge Distillation (KD) tries to preserve the knowledge capture in a previous version of the model by matching logits [26, 41], feature maps [11], or other information [18, 39, 45, 48, 51] in the new model. (ii) Exemplar Replay (ER) methods build a reservoir of samples or exemplars from old training rounds [3, 33, 40, 41, 43] and replay them in successive training phases as a way of recalling past knowledge. KD and ER are the starting point of our method.

Incremental object detection (IOD). IOD applies incremental learning to object detection specifically. This is more challenging than incremental image classification, as images can contain multiple objects, both of old and new types, with only the new types being annotated in any given training phase. Both KD and ER have been applied to detection before. [44] applies KD to the output of Faster R-CNN [13]. Inspired by this, recent IOD methods extended the KD framework to other detectors (*e.g.*, Faster-RCNN [42] and GFL [25]) by adding KD terms on the intermediate feature maps [12, 55, 58] and region proposal networks [7, 14, 37]. [19] proposes instead to store a set of exemplars and fine-tune the model on the exemplars after each incremental step. [30] proposes an adaptive sampling strategy to achieve more efficient exemplar selection for IOD.

However, existing IOD methods are designed based on conventional detectors such as Faster-RCNN [42] and GFL [25]. In this work, we show that a direct application of KD and ER to current state-of-the-art transformer-based detectors such as Deformable DETR [59] and UP-DETR [9] does not work well and we propose fixes to this issue.

Transformer-based object detection. DETection TRansformer (DETR) [4] proposes an elegant architecture for object detection based on a visual transformer [49]. Compared to pre-transformer approaches, DETR eliminates the need for non-maximum suppression in post-processing because self-attention can learn to remove duplicated detection by itself. This is achieved by using the Hungarian loss, matching each object hypothesis to exactly one target or background

using bipartite matching [47]. Deformable DETR [59] improves the performance of DETR, particularly for small objects, via sparse attention on multi-level feature maps. UP-DETR [9] leverages unsupervised learning to pre-train the parameters of the encoder and decoder in DETR to further boost the performance.

Our method does not fundamentally change these detectors and is in fact applicable to all similar ones. Instead, it proposes broadly-applicable changes that make transformer-based detectors work well in combination with KD and ER for the IOD problem.

3. Methodology

After defining the incremental detection problem (Sec. 3.1) and providing the necessary background (Sec. 3.2), we introduce ContinuaL DETection TRansformer (CL-DETR), a new method for incremental object detection that extends DETR-like detectors with knowledge distillation (KD; Sec. 3.3) and exemplar replay (ER; Sec. 3.4).

3.1. Incremental object detection

In *incremental object detection* (IOD) the goal is to train a detector in phases, where in each phase the model is only given annotations for a subset of the object categories. Formally, let $\mathcal{D} = \{(x, y)\}$ be a dataset of images x with corresponding object annotations y , such as COCO 2017 [27], and let $\mathcal{C} = \{1, \dots, C\}$ be the set of object categories. We adapt such a dataset for benchmarking IOD as follows. First, we partition \mathcal{D} and \mathcal{C} into M subsets $\mathcal{D} = \mathcal{D}_1 \cup \dots \cup \mathcal{D}_M$ and $\mathcal{C} = \mathcal{C}_1 \cup \dots \cup \mathcal{C}_M$, one for each training phase. For each phase i , we modify the samples $(x, y) \in \mathcal{D}_i$ so that y only contains annotations for objects of class \mathcal{C}_i and drop the others.²

In phase i of training, the model is only allowed to observe images \mathcal{D}_i with annotations for objects of types $\mathcal{C}_i \subset \mathcal{C}$. Notably, images can and do contain objects of any possible type \mathcal{C} , but only types \mathcal{C}_i are annotated in this phase. After phase i is complete, training switches to the next phase $i + 1$, so the model observes different images \mathcal{D}_{i+1} and annotations for objects of different types \mathcal{C}_{i+1} .

For exemplar replay, we relax this training protocol and allow the model to memorise a small number of *exemplars* $\mathcal{E}_i \subset \mathcal{D}_i$ from the previous phases. In this case, the model is trained on the union $\mathcal{D}_i \cup \mathcal{E}_{1:i-1}$ where $\mathcal{E}_{1:i-1} = \mathcal{E}_1 \cup \dots \cup \mathcal{E}_{i-1}$ forms the exemplar memory.

Note that this is a stricter and improved protocol compared to prior works in IOD [12, 44]. In these works, the model is still presented a subset of annotations restricted to classes \mathcal{C}_i in each phase; however, $\mathcal{D}_i \subset \mathcal{D}$ is defined as the subset of all images that contain objects of type \mathcal{C}_i . Because

²In this way, some images end up containing no annotated objects.

images contain a mix of object categories that can span different subsets \mathcal{C}_i , this means that different subsets \mathcal{D}_i can overlap, so that the same images can be observed multiple times in different phases. This violates the standard definition of incremental learning [16, 33, 41] which assumes that different samples are observed in different phases. Our setting retains this property.

3.2. Transformer-based detectors

State-of-the-art methods like DETR [4, 9, 28, 47, 56, 59] build on powerful visual transformers to solve the object detection problem. In order to motivate and explain our method, we first review briefly how they work.

With reference to Fig. 2, the model Φ takes as input an image $x \in \mathbb{R}^{3 \times H \times W}$ and outputs the object predictions $\hat{y} = \Phi(x)$ using a number of attention and self-attention layers. The output $\hat{y} = (\hat{y}_j)_{j \in \mathcal{N}}$ is a sequence $\mathcal{N} = \{1, \dots, N\}$ of object predictions $\hat{y}_j = (\hat{p}_j, \hat{b}_j)$, consisting of a class probability vector $\hat{p}_j : \mathcal{C} \cup \{\phi\} \rightarrow [0, 1]$ and a vector $b_j \in [0, 1]^4$ specifying the centre and size of the object bounding box relative to the image size. Note that the support of \hat{p}_j includes element ϕ that denotes the background class, or ‘no object’ (hence, \hat{p}_j has $C + 1$ dimensions).

The object predictions correspond to a fixed set of *object queries* internal to the model. Each query is thus mapped to an object instance or background. The order of the queries is conceptually immaterial, but queries are fixed and non-interchangeable after training. For instance, \hat{y}_1 is always the prediction that corresponds to the first query in the model. This is relevant for the application of KD.

For supervised training, the model is given ground truth object annotations $y = ((p_j, b_j))_{j \in \mathcal{N}}$ where p_j is the indicator vector of the category of the object and $b_j \in [0, 1]^4$ is its bounding box. Images usually contain fewer objects than the number N of hypotheses, so y is padded with background detections for which $p_i(\phi) = 1$ and b_i is arbitrary. The model is trained end-to-end to optimise the loss,

$$\mathcal{L}_{\text{DETR}}(\hat{y}, y) = \sum_{i \in \mathcal{N}} \langle -\log \hat{p}_{\hat{\sigma}_i}, p_i \rangle + \mathbf{1}_{c(p_i) \neq \phi} \mathcal{L}_{\text{box}}(\hat{b}_{\hat{\sigma}_i}, b_i), \quad (1)$$

where $c(p_i) = \arg\max_{c \in \mathcal{C} \cup \{\phi\}} p_i(c)$ is the class encoded by p_i , $\mathcal{L}_{\text{box}}(\hat{b}_{\hat{\sigma}_i}, b_i) = \gamma_1 \mathcal{L}_{\text{IoU}}(\hat{b}_{\hat{\sigma}_i}, b_i) + \gamma_2 \|\hat{b}_{\hat{\sigma}_i} - b_i\|_1$ is the bounding box prediction loss and $\hat{\sigma}$ is the best association of ground truth labels to object hypotheses, obtained by solving the matching problem,

$$\hat{\sigma} = \arg\max_{\sigma \in \mathcal{S}_{\mathcal{N}}} \sum_{i \in \mathcal{N}} \mathbf{1}_{c(p_i) \neq \phi} \left\{ -\langle \hat{p}_{\sigma_i}, p_i \rangle + \mathcal{L}_{\text{box}}(\hat{b}_{\sigma_i}, b_i) \right\} \quad (2)$$

using the Hungarian algorithm [22, 46] (see [4] for details).

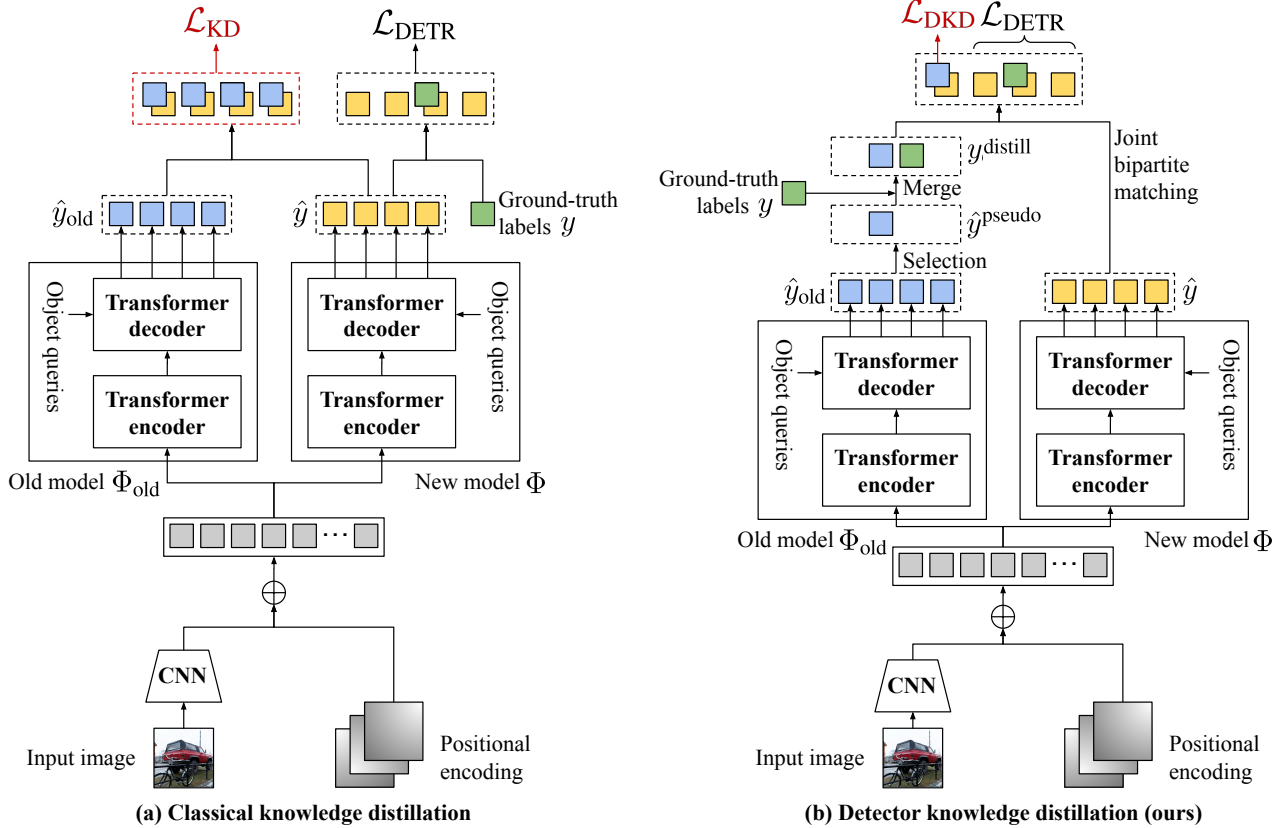


Figure 2. **(a) Classical knowledge distillation.** There are two issues when directly applying KD [15, 26] to the transformer-based detectors [4, 9, 59]. (i) Transformer-based detectors work by testing a large number of object hypotheses in parallel. Because the number of hypotheses is much larger than the typical number of objects in an image, most of them are negative, resulting in an unbalanced KD loss. (ii) Because both old and new object categories can co-exist in any given training image, the KD loss and regular training objective can provide contradictory evidence. **(b) Detector knowledge distillation (ours).** We select the most confident foreground predictions from the old model and use them as pseudo labels. We purposefully ignore background predictions because they are imbalanced and they can contradict the labels of the new classes available in the current phase. Then, we *merge* the pseudo labels for the old categories with the ground-truth labels for the new categories and use bipartite matching to train the model on the joint labels. This inherits the good properties of the original formulation such as ensuring one-to-one matching between labels and hypotheses and avoiding duplicate detections.

3.3. Detector knowledge distillation

In a multi-phase learning scenario, at the beginning of a new phase, the model is initialized as $\Phi \leftarrow \Phi^{old}$ where Φ^{old} is the model trained in the phase before. As the new data for the current phase is received, training the model Φ as normal by minimising eq. (1) leads to forgetting.

KD [15, 26] reduces forgetting by maintaining a copy of the old model and making sure that the outputs of the new and old models stay close. Applied to our transformer-based detectors, given a new training image-label pair (x, y) , one computes the old model's output $\hat{y}^{old} = \Phi^{old}(x)$ and, minimizes the sum of the $\mathcal{L}_{DETR}(\hat{y}, y)$ loss with the *knowledge distillation loss*

$$\mathcal{L}_{KD}(\hat{y}, \hat{y}^{old}) = \sum_{j \in \mathcal{N}} \left[\sum_{c \in \mathcal{C}} -\hat{p}_j(c) \log \hat{p}_j^{old}(c) \right] + \mathcal{L}_{box}(\hat{b}_j, \hat{b}_j^{old}).$$

This loss compares the output tokens of the new and old models, which makes sense since they depend on the same object queries, at least initially, and are thus in correspondence. However, we find that this loss is dominated by background information because most of the tokens predict background. Furthermore, transformer-based detectors aim to find one-to-one matchings between predictions and ground-truth labels without duplicates, which is not accounted for by the classical KD loss.

The key issue is that summing losses $\mathcal{L}_{DETR} + \mathcal{L}_{KD}$ as in standard KD fails to properly account for the *structure* of the labels, which is crucial for detection problems, particularly in an incremental learning setting. Specifically, the old model knows about all categories seen so far during training *except* the new categories that are annotated in the current phase. However, the new training images contain multiple objects, including the old types, which are thus *not* anno-

tated in the current phase. This means that $\mathcal{L}_{\text{DETR}}$ and \mathcal{L}_{KD} provide potentially contradictory supervision.

We thus suggest that, in a detection context, new and old knowledge should be fused in a *structured manner*. As illustrated in Fig. 2, we do so by selecting the most confident foreground predictions from the old model and using them as pseudo labels. We purposefully ignore background predictions because they are imbalanced and they can contradict the labels of the new classes available in the current phase. Then, we *merge* the pseudo labels for the old categories with the ground-truth labels for the new categories and use bipartite matching to train the model on the joint labels. This inherits the good properties of the original formulation such as ensuring one-to-one matching between labels and hypotheses and avoiding duplicate detections.

Formally, given the predictions \hat{y}^{old} from the old model, we first identify the subset $\mathcal{F} \subset \mathcal{N}$ of the ones that are predicted as foreground: $\mathcal{F} = \{j \in \mathcal{N} : \forall c \in \mathcal{C} : \hat{p}_j^{\text{old}}(c) > \hat{p}_j^{\text{old}}(\phi)\}$. Of these, we pick the subset $\mathcal{P} \subset \mathcal{F}$, $|\mathcal{P}| = K$ formed by the K most confident predictions, *i.e.*,

$$\forall i \in \mathcal{P}, j \in \mathcal{F} - \mathcal{P} : \max_{c \in \mathcal{C}} \hat{p}_i^{\text{old}}(c) > \max_{c \in \mathcal{C}} \hat{p}_j^{\text{old}}(c).$$

Finally, we further restrict the predictions to the subset $\mathcal{Q} \subset \mathcal{P}$ that does not overlap too much with the ground-truth labels for the new categories:

$$\mathcal{Q} = \{j \in \mathcal{P} : \forall i \in \mathcal{N} : c(p_i) \neq \phi \Rightarrow \text{IoU}(\hat{b}_j^{\text{old}}, b_i) \leq \lambda\}.$$

In the experiments, we set $\lambda = 0.7$. With remain with a filtered set of pseudo-labels:

$$\hat{y}^{\text{pseudo}} = (\hat{y}_j^{\text{old}})_{j \in \mathcal{Q}}. \quad (3)$$

Next, we distill knowledge from the current labels y and the pseudo-labels obtained from the old model into a single, coherent set of labels

$$y^{\text{distill}} = (y_i)_{i:c(p_i) \neq \phi} \oplus \hat{y}^{\text{pseudo}} \oplus y^{\text{bg}}, \quad (4)$$

where we concatenate the object labels for the new categories, the pseudo-labels, and enough background labels y^{bg} to pad y^{distill} to contain N elements.

In this manner, the distillation occurs at the level of the labels. The model is still trained by using eq. (1) as before, resulting in the *detector knowledge distillation* (DKD) loss:

$$\mathcal{L}_{\text{DKD}}(\hat{y}, y^{\text{distill}}) = \mathcal{L}_{\text{DETR}}(\hat{y}, y^{\text{distill}}). \quad (5)$$

Besides the usage of the distilled labels, the main difference between eqs. (1) and (5) is that, while the class distribution p_i for the new label is deterministic, it is *not* for the pseudo-labels. Plugged in eq. (1), this results in the standard distillation effect for categorical distributions trained using the cross entropy loss.

Algorithm 1: CL-DETR (the i -th phase)

Input: new category data \mathcal{D}_i ; old category exemplars $\mathcal{E}_{1:i-1}$; old model Φ^{old} .
Output: new model Φ ; exemplars $\mathcal{E}_{1:i}$.

- 1 Get \mathcal{D}_i and load $\mathcal{E}_{1:i-1}$ from memory;
- 2 Let $\Phi \leftarrow \Phi^{\text{old}}$;
- 3 **for** epochs **do**
- 4 **for** mini-batches $(x, y) \in \mathcal{D}_i \cup \mathcal{E}_{1:i-1}$ **do**
- 5 Let $\hat{y}^{\text{old}} \leftarrow \Phi^{\text{old}}(x)$;
- 6 Get \hat{y}^{pseudo} from \hat{y}^{old} and y using eq. (3);
- 7 Get y^{distill} from \hat{y}^{pseudo} and y using eq. (4);
- 8 Let $\hat{y} \leftarrow \Phi(x)$;
- 9 Get $\hat{\sigma}$ by matching y^{distill} to \hat{y} using eq. (2);
- 10 Compute $\mathcal{L}_{\text{DKD}}(\hat{y}, y^{\text{distill}})$ using eq. (5);
- 11 Update Φ via a gradient step.
- 12 Build the exemplar set $\mathcal{E}_{1:i}$ using Algorithm 2;
- 13 **for** epochs **do**
- 14 **for** mini-batches $(x, y) \in \mathcal{E}_{1:i}$ **do**
- 15 Let $\hat{y} \leftarrow \Phi(x)$;
- 16 Compute $\mathcal{L}_{\text{DETR}}(\hat{y}, y)$ using eq. (1);
- 17 Update Φ via a gradient step;
- 18 Save $\mathcal{E}_{1:i}$ to the memory.

Algorithm 2: Exemplar selection (the i -th phase)

Input: new category data \mathcal{D}_i ; old category exemplars $\mathcal{E}_{1:i-1}$; target number of exemplars R_i .
Output: exemplars $\mathcal{E}_{1:i}$.

- 1 Let $\mathcal{E}_i \leftarrow \{\}$;
- 2 **repeat**
- 3 Select $e \in \mathcal{D}_i$ according to eq. (6);
- 4 Let $\mathcal{E}_i \leftarrow \mathcal{E}_i \cup \{e\}$;
- 5 **until** R_i times;
- 6 Let $\mathcal{E}_{1:i} \leftarrow \mathcal{E}_i \cup \mathcal{E}_{1:i-1}$.

3.4. Distribution-preserving calibration

ER methods, which store a small number of exemplars and replay them in future phases, are shown to be effective in preserving the old category knowledge in IOD [19, 30], but can suffer from the severe imbalance between old and new category annotations. Incremental learning methods for classification [16, 31, 53] usually use re-balancing strategies to address the imbalance problem. They create a category-balanced subset of the data and finetune some model components (*e.g.*, the classifier) on it. However, such strategies do not apply directly to the IOD setting. First, the class distribution in detection is far from balanced, and a better strategy is to match the natural data distribution instead of the uniform one. Second, because there are multi-

Setting	Method	Detection baseline	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
70+10	ERD [12]	UP-DETR	36.2 \pm 0.3	54.8 \pm 0.4	39.3 \pm 0.4	20.8 \pm 0.3	39.3 \pm 0.5	47.9 \pm 0.3
	CL-DETR (ours)	UP-DETR	37.6 \pm 0.2	56.5 \pm 0.4	39.4 \pm 0.3	20.5 \pm 0.3	39.1 \pm 0.4	49.9 \pm 0.3
	LwF [26]	Deformable DETR	24.5 \pm 0.3	36.6 \pm 0.2	26.7 \pm 0.4	12.4 \pm 0.2	28.2 \pm 0.4	35.2 \pm 0.4
	iCaRL [41]	Deformable DETR	35.9 \pm 0.4	52.5 \pm 0.3	39.2 \pm 0.3	19.1 \pm 0.3	39.4 \pm 0.5	48.6 \pm 0.3
	ERD [12]	Deformable DETR	36.9 \pm 0.4	55.7 \pm 0.4	40.1 \pm 0.4	21.4 \pm 0.3	39.6 \pm 0.3	48.7 \pm 0.3
	CL-DETR (ours)	Deformable DETR	40.1\pm0.3	57.8\pm0.4	43.7\pm0.3	23.2\pm0.3	43.2\pm0.2	52.1\pm0.3
40+40	ERD [12]	UP-DETR	35.4 \pm 0.4	55.1 \pm 0.3	38.3 \pm 0.3	17.9 \pm 0.4	39.0 \pm 0.3	49.8 \pm 0.3
	CL-DETR (ours)	UP-DETR	37.0 \pm 0.2	56.2\pm0.2	39.1 \pm 0.4	20.9\pm0.2	38.9 \pm 0.3	49.2 \pm 0.3
	LwF [26]	Deformable DETR	23.9 \pm 0.2	41.5 \pm 0.3	25.0 \pm 0.3	12.0 \pm 0.4	26.4 \pm 0.3	33.0 \pm 0.5
	iCaRL [41]	Deformable DETR	33.4 \pm 0.4	52.0 \pm 0.3	36.0 \pm 0.2	18.0 \pm 0.3	36.4 \pm 0.3	45.5 \pm 0.4
	ERD [12]	Deformable DETR	36.0 \pm 0.2	55.2 \pm 0.2	38.7 \pm 0.3	19.5 \pm 0.2	38.7 \pm 0.3	49.0 \pm 0.4
	CL-DETR (ours)	Deformable DETR	37.5\pm0.3	55.1 \pm 0.4	40.3\pm0.2	20.9\pm0.2	40.8\pm0.4	50.7\pm0.2

Table 1. IOD results (%) on COCO 2017. In the $A + B$ setup, in the first phase, we observe a fraction $\frac{A}{A+B}$ of the training samples with A categories annotated. Then, in the second phase, we observe the remaining $\frac{B}{A+B}$ of the training samples, where B new categories are annotated. We test settings $A + B = 40 + 40$ and $70 + 10$. Exemplar replay is applied for all methods except for LwF [26]. We run experiments for three different categories and data orders and report the average AP with 95% confidence interval.

ple objects in each image, it is non-trivial to create a subset of exemplar images with a set number of objects for each category. We address these issues next.

Selecting exemplars to match the training distribution.

Called during phase i , Algorithm 2 produces a new exemplar subset \mathcal{E}_i whose distribution matches as well as possible the distribution of categories in the subset \mathcal{D}_i of the data. This is achieved by adding to \mathcal{E}_i a set number R_i of one exemplar $e^* \in \mathcal{D}_i$, one at a time, chosen by minimizing the Kullback-Leibler divergence [23] between the category marginals of \mathcal{E}_i and \mathcal{D}_i :

$$e^* \leftarrow \sum_{c \in \mathcal{C}_i} p_{\mathcal{D}_i}(c) \log p_{\mathcal{E}_i \cup \{e\}}(c), \quad (6)$$

where $p_{\mathcal{D}}(c)$ denotes the probability of category c in dataset \mathcal{D} . Then, the overall exemplar set $\mathcal{E}_{1:i} = \mathcal{E}_i \cup \mathcal{E}_{1:i-1}$ is obtained as the union of the new subset just found and the previous exemplar set $\mathcal{E}_{1:i-1}$. Because classes in different subsets \mathcal{D}_i are disjoint, this also means that, by the end of the training, the distribution of classes in $\mathcal{E}_{1:M}$ approximates the one of the overall training set \mathcal{D} .

Learning using balanced data. In order to use the available data as well as possible while balancing the detector Φ , in each phase we update it in two steps. In the first step, the model is trained using the DKD loss on all the available data $\mathcal{D}_i \cup \mathcal{E}_{1:i-1}$ given by the union of the current data subset \mathcal{D}_i and the exemplar memory $\mathcal{E}_{1:i-1}$ carried over the previous training phases. In the second step, the model is fine-tuned using the new exemplar set $\mathcal{E}_{1:i}$, ignoring \mathcal{D}_i and using only the DETR loss, using fewer data but achieving better calibration. The overall algorithm is given in Algorithm 1.

4. Experiments

We evaluate CL-DETR on COCO 2017 using two transformer-based detectors, Deformable DETR and UP-DETR [9, 59] as the baselines and achieve consistent improvements compared to the baselines and a direct application of KD and ER. Below we describe the dataset and implementation details (Sec. 4.1) followed by results and analyses (Sec. 4.2).

4.1. Dataset and implementation details

Dataset and evaluation metrics. We conduct IOD experiments on COCO 2017 [27], which is widely used in related works [9, 12, 38, 59]. Following [12], the standard COCO metrics are used for evaluation, *i.e.*, AP , AP_{50} , AP_{75} , AP_S , AP_M , and AP_L . In the ablation study, we introduce a new metric, *forgetting percentage points* (FPP), measuring the difference between the AP of the first and last phase models on the categories observed in the first phase.

Experiment setup. We conduct IOD experiments in the following setting. **Two-phase setting:** In the $A + B$ setup, in the first phase, we observe a fraction $\frac{A}{A+B}$ of the training samples with A categories annotated. Then, in the second phase, we observe the remaining $\frac{B}{A+B}$ of the training samples, where B new categories are annotated. We test settings $A+B = 40+40$ and $70+10$. **Multiple-phase setting:** In the $40+X \times Y$ setup, in the first phase, we observe half of the training samples with 40 categories annotated. In each following phase, we observe $\frac{1}{2Y}$ of the training samples we have never seen before with annotations for X new categories. We run experiments for $40+20 \times 2$ and $40+10 \times 4$. We repeat each experiment three times, randomizing the order of categories and data in the different phases, and report the average APs. The total memory budget for the exem-




Figure 3. IOD results (AP/AP₅₀, %) on COCO 2017 in the 40 + 20 × 2 and 40 + 10 × 4 settings. Our method is based on Deformable DETR. Comparing methods: Upper Bound (the results of joint training with all previous data accessible in each phase), ERD [12], SID [38], and RILOD [24]. The results of the related works are from [12]. We use the same data split as [12] for a fair comparison.

Row	Knowledge distillation (KD)	Joint bipartite matching	Pseudo label selection	Exemplar replay (ER)	Distribution preserving calibration	All categories ↑				Old categories ↑				FPP ↓			
						AP	AP _S	AP _M	AP _L	AP	AP _S	AP _M	AP _L	AP	AP _S	AP _M	AP _L
1						4.2	1.6	4.7	5.8	0.7	0.2	0.8	0.8	42.6	25.6	45.1	56.7
2	✓					24.5	12.4	28.2	35.2	24.0	12.3	27.7	34.4	19.3	13.5	18.2	23.1
3	✓	✓				30.3	19.5	33.0	39.0	33.4	21.8	36.4	43.2	9.9	4.0	9.5	14.3
4	✓	✓	✓			33.9	16.3	37.1	49.2	33.9	16.6	36.8	50.0	9.4	9.2	9.1	7.5
5	✓	✓	✓	✓		37.9	20.8	40.9	50.4	39.0	21.6	41.7	52.3	4.3	4.2	4.2	5.2
6	✓	✓	✓		✓	40.1	23.2	43.2	52.1	41.8	24.5	44.7	54.6	1.5	1.3	1.2	2.9

Table 2. Ablation results (%) for KD and ER, using Deformable DETR [59] on COCO 2017 in the 70 + 10 setting. “All categories” (higher is better) denote the results of the last phase model on 80 categories. “Old categories” (higher is better) denote the results of the last phase model on 70 categories observed in the first phase. “Forgetting percentage points (FPP)” (lower is better) show the difference between the AP of the first-phase model and the last-phase model on 70 categories observed in the first phase. The baseline (row 1) is finetuning the model without IOD techniques. Our method (CL-DETR) is shown in row 6.

plars is set as 10% of the total dataset size.

Implementation details. We follow [9, 59] and use an ImageNet pre-trained ResNet-50 backbone. For the experiments on Deformable DETR [59], we use the standard configurations without their iterative bounding box refinement mechanism and the two-stage Deformable DETR. We train the model for 50 (Deformable DETR) and 150 epochs (UP-DETR), following the original implementations [9, 59]. In order to apply our distribution-preserving calibration (Sec. 3.4), we train the coarse Deformable DETR (UP-DETR) model for 40 (120) epochs and perform calibration for 10 (30) epochs to preserve the total number of epochs.

4.2. Results and analyses

Two-phase setting. Table 1 shows that, in the two-phase settings 70 + 10 and 40 + 40, applying CL-DETR to Deformable DETR [59] and UP-DETR [9] consistently performs better than the state-of-the-art [12] and other IOD methods [26, 41]. In particular, Deformable DETR [59] w/ ours achieves the highest AP, *e.g.*, 40.1% and 37.5% in the 70 + 10 and 40 + 40 settings, respectively. The performance

gap is larger with more categories in the 1-st phase. *E.g.*, the AP differences between our method and [12] are 3.2 and 1.5 percentage points when we observe 70 and 40 categories in the first phase, respectively, likely due to CL-DETR benefiting more from a well-pre-trained model.

Multiple-phase setting. Figure 3 evaluates CL-DETR in the multiple-phase setting with large gains compared to other IOD methods in both the 40 + 20 × 2 and 40 + 10 × 4 experimental variants. The relative advantage of CL-DETR increases with the number of phases. For instance, our method improves the AP of [12] by 2.9 percentage points in the 40 + 20 × 2 setting and by 7.4 percentage points in the 40 + 10 × 4 setting. This suggests that the advantage of CL-DETR shows more in challenging settings, where the forgetting problem is stronger due to the larger number of training phases.

Ablation study for DKD. In Tab. 2 (Rows 1–4) we ablate our DKD approach. By comparing row 2 to row 1, we observe that classical KD significantly improves the IOD performance compared to the baseline (*i.e.*, finetuning the model without IOD techniques), but still results in large

Row	Setting	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
1	$K=5$	39.7	57.4	43.1	22.7	42.6	52.7
2	$K=10$	40.1	57.8	43.7	23.2	43.2	52.1
3	$K=20$	39.9	57.8	43.2	23.5	42.9	51.7
4	$p \geq 0.1$	39.3	57.1	42.9	22.6	42.3	52.5
5	$p \geq 0.3$	39.6	57.5	43.0	23.2	42.4	52.2
6	$p \geq 0.5$	39.2	56.8	42.4	22.3	41.9	51.8

Table 3. Ablation result (%) for different pseudo label selection strategies on COCO 2017 using the 70 + 10 setting. Rows 1–3 show the results for using different K when selecting top- K most-confident non-background predictions. Rows 4–6 show the results for using different thresholds p of the prediction scores to select the non-background predictions.

overall forgetting: 19.3 FPP. Comparing row 3 to row 2, we can see that joint bipartite matching works well and boosts the AP of all categories by 5.8 percentage points compared to conventional KD. The reason is that joint bipartite matching helps ensure a one-to-one matching between objects and hypotheses and discourages duplicate detections. Comparing row 4 to row 3, our pseudo label selection further improves the AP and reduces forgetting, helping the model to ignore the redundant background information and reducing conflicts between old and new labels.

Ablation study for ER. In Tab. 2 (Rows 5–6), we ablate our ER method. Comparing row 6 to row 5, we can see that the calibration strategy of Sec. 3.4 boosts both the all-category and old-category performance, by 1.8 and 2.1 percentage points respectively, compared to using conventional ER [30, 41]. It also helps to overcome the catastrophic forgetting problem in IOD, reducing the AP forgetting by 2.1 percentage points. This is because the conventional ER balances the sample distributions, changing the category distribution of the training set, whereas our method preserves it, thus improving performance.

Ablation study for pseudo label selection strategies. In Tab. 3, we show the results for two pseudo-label selection strategies: (1) selecting top- K most-confident non-background predictions (Rows 1–3); and (2) selecting the predictions using a threshold for the prediction scores (Rows 4–6). We observe the first strategy works better, with peak AP when $K=10$. The maximum performance difference is only 0.4 percentage points when using different values for K . This indicates our method is robust to its hyperparameter settings.

Visualizations. Figure 4 visualizes the old category pseudo (blue) and ground-truth (green) bounding boxes in some training samples in COCO 2017. In Fig. 4 (a, b), CL-DETR generates accurate pseudo bounding boxes that exactly match the ground-truth ones. This shows the effective-




Figure 4. Visualizations of the old category pseudo (blue) and ground-truth (green) bounding boxes on COCO 2017 using the 70 + 10 setting. (a, b): Our method generates accurate pseudo bounding boxes that exactly match the ground-truth ones. (c, d): When there are too many annotations in the images, generated pseudo bounding boxes cannot cover all ground-truth ones. However, the pseudo bounding boxes are still focused on the foreground objects.

ness of our pseudo-label selection strategy. In Fig. 4 (c, d), CL-DETR fails to generate pseudo bounding boxes for all objects in the images when there are too many. This is explained by our strategy of selecting the top- K most-confident non-background bounding boxes as the pseudo-labels followed by removing the ones that overlap with the new category ground-truth labels excessively. In this manner, the number of pseudo bounding boxes is always smaller than K . The trade-off, justified by our improvements in the experiments, is to prefer correct although possibly incomplete annotations to contradictory or noisy ones.

5. Conclusions

This paper introduced CL-DETR, a novel IOD method that can effectively use KD and ER in transformer-based detectors. CL-DETR improves the standard KD loss by introducing DKD which selects the most informative predictions from the old model, rejecting redundant background predictions, and ensuring that the distilled information is consistent with the new ground-truth evidence. CL-DETR also improves ER by selecting exemplars to match the distribution of the training set. CL-DETR is fairly generic and can be easily applied to different transformer-based detectors, including Deformable DETR [59] and UP-DETR [9], achieving large improvements. We have also defined a more realistic IOD benchmark protocol that avoids using duplicated images in different training phases. In the future, we plan to extend our method to more challenging settings such as online learning.

Ethics. We use the COCO dataset in a manner compatible with their terms; this data contains personal information (faces). For further details on ethics, data protection, and copyright please see <https://www.robots.ox.ac.uk/~vedaldi/research/union/ethics.html>.

Acknowledgments. C. R. and A. V. are supported by ERC-UNION-CoG-101001212. C. R. is also supported by VisualAI EP/T028572/1.

References

- [1] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *CVPR*, pages 3366–3375, 2017. [2](#)
- [2] Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *CVPR*, pages 11254–11263, 2019. [1](#), [2](#)
- [3] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *CVPR*, pages 8218–8227, 2021. [2](#)
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020. [1](#), [2](#), [3](#), [4](#)
- [5] Francisco M. Castro, Manuel J. Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, pages 241–257, 2018. [1](#)
- [6] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with A-GEM. In *ICLR*, 2019. [2](#)
- [7] Li Chen, Chunyan Yu, and Lvcai Chen. A new knowledge distillation for incremental object detection. In *IJCNN*, pages 1–7, 2019. [2](#)
- [8] Zhiyuan Chen and Bing Liu. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207, 2018. [2](#)
- [9] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. UP-DETR: Unsupervised pre-training for object detection with transformers. In *CVPR*, pages 1601–1610, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [10] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *TPAMI*, 44(7):3366–3385, 2021. [2](#)
- [11] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, pages 86–102, 2020. [1](#), [2](#)
- [12] Tao Feng, Mang Wang, and Hangjie Yuan. Overcoming catastrophic forgetting in incremental object detection via elastic response distillation. In *CVPR*, pages 9427–9436, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [13] Ross Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. [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 *ICME*, pages 1–6, 2019. [2](#)
- [15] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. In *NIPS Workshops*, 2014. [4](#)
- [16] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019. [1](#), [2](#), [3](#), [5](#)
- [17] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *CVPR*, pages 3957–3966, 2021. [1](#)
- [18] K. J. Joseph, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Vineeth N Balasubramanian. Energy-based latent aligner for incremental learning. In *CVPR*, pages 7452–7461, 2022. [2](#)
- [19] K. J. Joseph, Salman H. Khan, Fahad Shahbaz Khan, and Vineeth N. Balasubramanian. Towards open world object detection. In *CVPR*, pages 5830–5840, 2021. [2](#), [5](#)
- [20] K. J. Joseph, Jathushan Rajasegaran, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Incremental object detection via meta-learning. *TPAMI*, 2021. [1](#)
- [21] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *PNAS*, pages 3521–3526, 2017. [1](#)
- [22] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. [3](#)
- [23] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951. [6](#)
- [24] 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 Songqing Chen, Ryokichi Onishi, Ganesh Ananthanarayanan, and Qun Li, editors, *SEC*, pages 113–126, 2019. [7](#)
- [25] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *NeurIPS*, 2020. [2](#)
- [26] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2018. [1](#), [2](#), [4](#), [6](#), [7](#)
- [27] 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, 2014. [3](#), [6](#)
- [28] Fanfan Liu, Haoran Wei, Wenzhe Zhao, Guozhen Li, Jingquan Peng, and Zihao Li. WB-DETR: Transformer-based detector without backbone. In *ICCV*, pages 2979–2987, 2021. [3](#)
- [29] Xialei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, Andrew D Bagdanov, Shangling Jui, and Joost van de Weijer. Generative feature replay for class-incremental learning. In *CVPR Workshops*, pages 226–227, 2020. [1](#)
- [30] 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](#), [5](#), [8](#)

- [31] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, pages 2544–2553, 2021. 5
- [32] Yaoyao Liu, Bernt Schiele, and Qianru Sun. RMM: reinforced memory management for class-incremental learning. In *NeurIPS*, pages 3478–3490, 2021. 1
- [33] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *CVPR*, pages 12245–12254, 2020. 1, 2, 3
- [34] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *NIPS*, pages 6467–6476, 2017. 2
- [35] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*, volume 24, pages 109–165. Elsevier, 1989. 1
- [36] K. McRae and P. Hetherington. Catastrophic interference is eliminated in pre-trained networks. In *CogSci*, 1993. 1
- [37] Can Peng, Kun Zhao, and Brian C. Lovell. Faster ILOD: incremental learning for object detectors based on faster RCNN. *Pattern Recognition Letter*, 140:109–115, 2020. 1, 2
- [38] 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. *CVIU*, 210:103229, 2021. 6, 7
- [39] Mozghan PourKeshavarzi, Guoying Zhao, and Mohammad Sabokrou. Looking back on learned experiences for class/task incremental learning. In *ICLR*, 2022. 2
- [40] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. GDumb: A simple approach that questions our progress in continual learning. In *ECCV*, pages 524–540, 2020. 2
- [41] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, pages 5533–5542, 2017. 1, 2, 3, 6, 7, 8
- [42] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *TPAMI*, 39(6):1137–1149, 2017. 2
- [43] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NIPS*, pages 2990–2999, 2017. 2
- [44] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. Incremental learning of object detectors without catastrophic forgetting. In *ICCV*, pages 3420–3429, 2017. 1, 2, 3
- [45] Christian Simon, Piotr Koniusz, and Mehrtash Harandi. On learning the geodesic path for incremental learning. In *CVPR*, pages 1591–1600, 2021. 2
- [46] Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *CVPR*, pages 2325–2333, 2016. 3
- [47] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris M Kitani. Rethinking transformer-based set prediction for object detection. In *ICCV*, pages 3611–3620, 2021. 3
- [48] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, Xing Wei, and Yihong Gong. Topology-preserving class-incremental learning. In *ECCV*, pages 254–270, 2020. 2
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NIPS*, pages 5998–6008, 2017. 2
- [50] Eli Verwimp, Kuo Yang, Sarah Parisot, Hong Lanqing, Steven McDonagh, Eduardo Pérez-Pellitero, Matthias De Lange, and Tinne Tuytelaars. Re-examining distillation for continual object detection. In *BMVC*, 2022. 1
- [51] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *ECCV*, 2022. 2
- [52] Liyuan Wang, Xingxing Zhang, Kuo Yang, Longhui Yu, Chongxuan Li, Lanqing Hong, Shifeng Zhang, Zhenguo Li, Yi Zhong, and Jun Zhu. Memory replay with data compression for continual learning. In *ICLR*, 2022. 1
- [53] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, pages 374–382, 2019. 5
- [54] Binbin Yang, Xincheng Deng, Han Shi, Changlin Li, Gengwei Zhang, Hang Xu, Shen Zhao, Liang Lin, and Xiaodan Liang. Continual object detection via prototypical task correlation guided gating mechanism. In *CVPR*, pages 9255–9264, 2022. 1
- [55] Dongbao Yang, Yu Zhou, Aoting Zhang, Xurui Sun, Dayan Wu, Weiping Wang, and Qixiang Ye. Multi-view correlation distillation for incremental object detection. *Pattern Recognition*, 131:108863, 2022. 2
- [56] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. DINO: DETR with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. 3
- [57] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shutao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, pages 13208–13217, 2020. 1
- [58] Wang Zhou, Shiyu Chang, Norma Sosa, Hendrik Hamann, and David Cox. Lifelong object detection. *arXiv preprint arXiv:2009.01129*, 2020. 2
- [59] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: deformable transformers for end-to-end object detection. In *ICLR*, 2021. 1, 2, 3, 4, 6, 7, 8