

YOLO-IOD: Towards Real Time Incremental Object Detection

Shizhou Zhang¹, Xueqiang Lv¹, Yinghui Xing^{1*}, Qirui Wu¹, Di Xu², Chen Zhao¹, Yanning Zhang¹

¹Northwestern Polytechnical University, China
² Huawei, China

Abstract

Current methods for incremental object detection (IOD) primarily rely on Faster R-CNN or DETR series detectors; however, these approaches do not accommodate the real-time YOLO detection frameworks. In this paper, we first identify three primary types of knowledge conflicts that contribute to catastrophic forgetting in YOLO-based incremental detectors: foreground-background confusion, parameter interference, and misaligned knowledge distillation. Subsequently, we introduce **YOLO-IOD**, a real-time Incremental Object Detection (IOD) framework that is constructed upon the pre-trained YOLO-World model, facilitating incremental learning via a stage-wise parameter-efficient fine-tuning process. Specifically, YOLO-IOD encompasses three principal components: 1) Conflict-Aware Pseudo-Label Refinement (CPR), which mitigates the foreground-background confusion by leveraging the confidence levels of pseudo labels and identifying potential objects relevant to future tasks. 2) Importance-based Kernel Selection (IKS), which identifies and updates the pivotal convolution kernels pertinent to the current task during the current learning stage. 3) Cross-Stage Asymmetric Knowledge Distillation (CAKD), which addresses the misaligned knowledge distillation conflict by transmitting the features of the student target detector through the detection heads of both the previous and current teacher detectors, thereby facilitating asymmetric distillation between existing and newly introduced categories. We further introduce **LoCo COCO**, a more realistic benchmark that eliminates data leakage across stages. Experiments on both conventional and LoCo COCO benchmarks show that YOLO-IOD achieves superior performance with minimal forgetting.

Introduction

The goal of Incremental Object Detection (IOD) is to consistently acquire new object classes over learning of successive tasks, while retaining the knowledge of those learned in the past. While recent works (Mo et al. 2024; Liu et al. 2023) have made significant progress in IOD, most existing methods are built upon detectors such as Faster R-CNN (Ren et al. 2015) or DETR (Carion et al. 2020). When these developed methods are applied to real-time YOLO series detectors (Redmon et al. 2016), they usually suffer significant

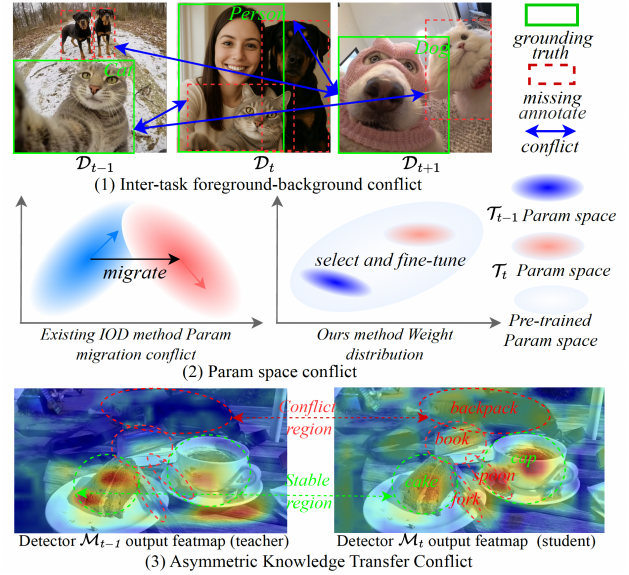


Figure 1: Illustration of the three major *knowledge conflicts*: (a) Foreground-background confusion, (b) Parameter interference and (c) Misaligned knowledge distillation conflict.

drops in generalization performance and struggle to preserve prior knowledge of previous classes, mainly due to knowledge conflicts across tasks.

In this work, we firstly identify the underlying causes of forgetting in YOLO-based incremental detectors as three major types of *knowledge conflicts* as shown in Fig. 1: 1) *Foreground-background confusion*, where unannotated objects of previous and future tasks are misclassified as background during training. This problem is critical for YOLO detectors, which rely on aggressive data augmentation techniques like Mosaic and MixUp that assume accurate annotations. In IOD settings, noise from pseudo labels is amplified by these augmentations, compromising model performance. 2) *Parameter interference*. This issue arises because different tasks frequently depend on intersecting parameter subsets within the model. Updates for new tasks can alter these shared parameters, thereby potentially disrupting previously acquired representations and leading to catastrophic forgetting of earlier tasks. 3) *Misaligned knowledge distil-*

*Corresponding Author: xyh_7491@nwpu.edu.cn
 Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

lation conflict, where the teacher and student models are optimized for mismatched class distributions, violating the core assumption in standard knowledge distillation that both models share consistent learning objectives. YOLO-series detectors, known for making dense predictions across spatial grids, are significantly affected by this problem.

To overcome above challenges, we propose YOLO-IOD, a real-time IOD framework that builds upon the pre-trained YOLO-World (Cheng et al. 2024) model and performs parameter efficient fine-tuning during each incremental learning stage. YOLO-IOD comprises three main modules designed to address the identified knowledge conflicts: 1) Conflict-Aware Pseudo-Label Refinement (CPR), which alleviates foreground-background confusion by incorporating two strategies: Enhanced Pseudo-label Loss, which improves the reliability of supervision from pseudo-labels by weighting predictions based on their confidence and uncertainty and Clustered Unknown Pseudo Labeling, which identifies potential objects from future tasks via performing open-vocabulary object detection and feature-space clustering. 2) Importance-based Kernel Selection (IKS), which mitigates parameter interference by selecting and fine-tuning only important task-relevant convolution kernels based on Fisher-based differential importance estimation. and 3) Cross-Stage Asymmetric Knowledge Distillation (CAKD), which resolves misaligned knowledge distillation conflict by passing cross-stage student features through detection heads of both old and current teacher detectors, enabling asymmetric distillation across old and new categories.

Moreover, current IOD benchmarks are not tailored for practical applications as they divide categories arbitrarily, disregarding natural co-occurrence of classes, and permit images to repeat across incremental stages. In practical situations, certain categories such as cars and pedestrians frequently occur together, whereas others, like cars and boats, do not. Images overlapping between stages must be strictly prevented to avoid data leakage. This raises questions about whether the performance observed on current benchmarks can effectively translate to real-world applications, particularly for recent IOD methods that depend on pseudo labels. Therefore, we introduce LoCo COCO, a novel benchmark designed to eliminate inter-stage image overlap and adhere to category co-occurrence statistics, offering a more equitable and realistic evaluation benchmark for IOD.

Our contributions can be summarized as follows:

- We introduce YOLO-IOD, an integrated and real-time IOD framework, and pinpoint three causes of forgetting: foreground-background confusion, parameter interference, and misaligned knowledge distillation conflict.
- YOLO-IOD incorporates three innovative modules aimed at mitigating forgetting, with particular emphasis on the dual-teacher CAKD module. This module addresses the challenge of misaligned knowledge distillation by channeling the target student detector’s features through the detection heads of both the former and current teacher detectors.
- We introduce LoCo COCO, a practical benchmark designed to remove image overlap between stages and con-

sider category co-occurrence, allowing for a more equitable assessment of incremental object detection.

- Extensive experiments on conventional COCO and LoCo COCO benchmarks under multiple settings show that our method consistently achieves state-of-the-art performance while maintaining real-time inference speed.

Related Works

Incremental Learning

Incremental learning enables models to acquire new information without forgetting prior knowledge (Schlimmer and Granger Jr 1986; Li and Hoiem 2017). Approaches include *rehearsal-based methods*, which replay or generate samples from previous tasks (Rebuffi et al. 2017; Lopez-Paz and Ranzato 2017; Shin et al. 2017); *regularization-based methods*, which constrain important parameters to maintain stability (Kirkpatrick et al. 2017; Luo et al. 2025; Yang et al. 2023; Wu et al. 2025); *architecture-based methods*, which assign task-specific subnetworks (Von Oswald et al. 2019; Rypešć et al. 2024); and *knowledge distillation*, which transfers information from past tasks via teacher-student learning (Tao et al. 2020; Chen and Chang 2023; Asadi et al. 2023). Recently, *continual learning with foundation models* have also attracted many interests, including approaches such as L2P (Wang et al. 2022) CODA-Prompt (Smith et al. 2023) and VPT-NSP² (Lu et al. 2024).

Incremental Object Detection

Compared to classification, IOD involves more complex scenarios and greater challenges due to the requirement of both localizing and classifying the objects. In IOD, distillation is particularly effective due to the rich feature representations and multi-level supervision signals available in detection frameworks. ERD (Feng, Wang, and Yuan 2022) introduces elastic response distillation that adaptively transfers knowledge. BPF (Mo et al. 2024) bridges past-future knowledge through dual teacher models, CL-DETR (Liu et al. 2023) applies distillation within DETR architectures, emphasizing reliable predictions from old models. Nevertheless, these methodologies tackle the issue of knowledge distillation by selecting only those old task outputs that do not coincide with new labels to serve as the distillation targets. This strategy is unsuitable for YOLO-style dense predictions. Moreover, this approach can only distill partial knowledge from teacher models, failing to fully resolve the underlying conflict. We propose CAKD which provides cross-stage asymmetric knowledge distillation by suppressing the responses of unrelated features with the detection heads.

Preliminaries and Benchmark

Problem Definition

In IOD, tasks arrive sequentially in an ordered sequence: $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_t, \dots, \mathcal{T}_n\}$. Each task \mathcal{T}_t aims to learn a specific set of object categories \mathcal{C}_t , and $\mathcal{C}_i \cap \mathcal{C}_j = \emptyset, \forall i \neq j$. For each task \mathcal{T}_t , a dataset is provided: $\mathcal{D}_t = \{(\mathcal{X}_t^i, \mathcal{Y}_t^i)\}_{i=1}^{N_t}$, where \mathcal{X}_t^i denotes the i -th image and \mathcal{Y}_t^i is its corresponding annotation. Importantly, the dataset \mathcal{D}_t is

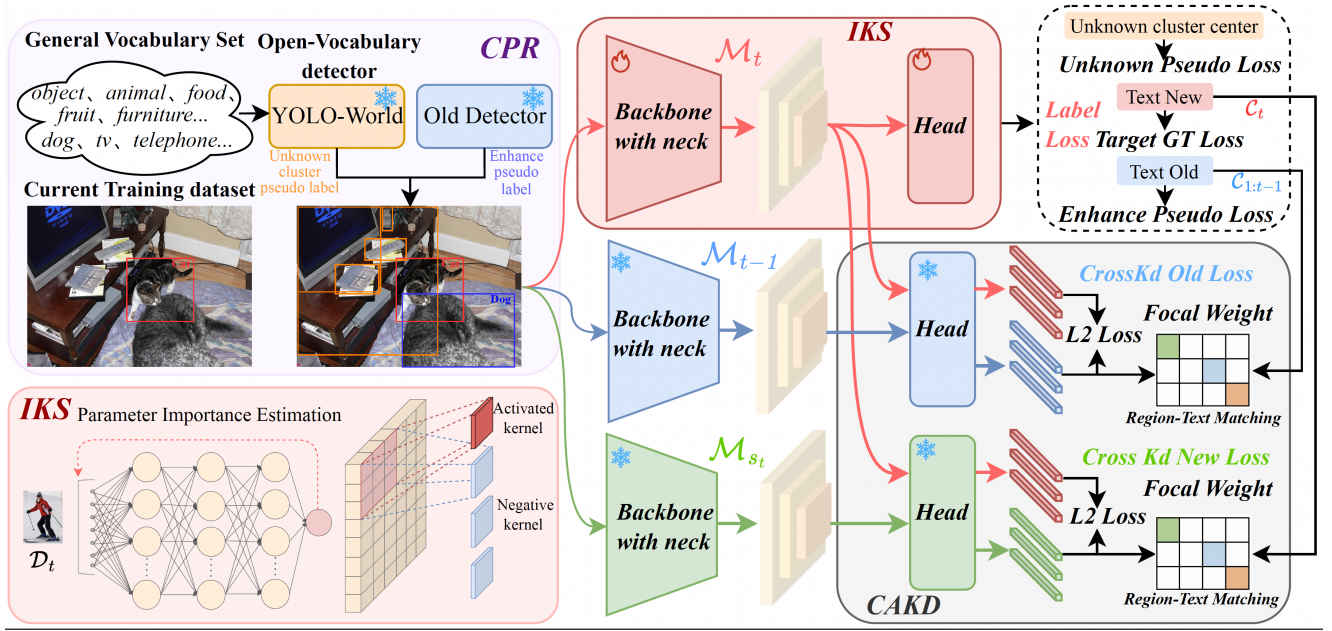


Figure 2: Overall architecture of YOLO-IOD. Our method integrates three components to address knowledge conflicts in IOD: 1) Conflict-aware Pseudo-label Refinement (CPR); 2) Importance-based Kernel Selection (IKS); and 3) Cross-Stage Asymmetric Knowledge Distillation (CAKD).

annotated only for categories in \mathcal{C}_t . That is, objects belonging to $\mathcal{C}_{1:t-1} \cup \mathcal{C}_{t+1:n}$ are not annotated, even if presented in the image. During the training of task \mathcal{T}_t , only \mathcal{D}_t is accessible. The goal of IOD is to train an object detector \mathcal{M}_t at each stage, such that it can correctly detect objects from both the current task’s classes \mathcal{C}_t and all seen classes $\mathcal{C}_{1:t-1}$.

YOLO-World

Open-vocabulary object detection (OVOD) enables detectors to recognize arbitrary categories guided by textual input. Recent methods such as GLIP (Li et al. 2022), Grounding DINO (Liu et al. 2024), and YOLO-World (Cheng et al. 2024) learn region-text alignment on large-scale vision-language data. Our method builds upon the pre-trained YOLO-World model. Images are processed by visual encoder f_v and category texts by text encoder f_t , with features fused using RepVL-PAN: $\mathbf{V}, \mathbf{P} = f_{\text{RepVL}}(f_v(I), f_t(T))$ where $\mathbf{V} = \{\mathbf{e}_k\}_{k=1}^K$ are region-level visual embeddings and $\mathbf{P} = \{\mathbf{p}_j\}_{j=1}^C$ are text prototypes. Classification scores are computed as $s_{k,j} = \eta \cdot \langle \text{Norm}(\mathbf{e}_k), \text{Norm}(\mathbf{p}_j) \rangle + \zeta$.

LoCo COCO Benchmark

In prior IOD benchmark, each stage t typically selects all images that contain objects of categories \mathcal{C}_t from the full dataset. Real-world images frequently contain objects from various categories, meaning that a single image will be used across several training stages. Based on the statistics, each image appears in an average of 1.84 stages in the 20+20 4 stage setting. This overlap challenges the foundational premise of continual learning and artificially inflates the effectiveness of pseudo-labeling methods by allowing detec-

tors to generate pseudo-labels on reused training images.

To address this issue, we introduce a new data partitioning protocol named as Low Co-occurrence COCO (LoCo COCO). We first construct a category co-occurrence matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, where \mathbf{A}_{ij} denotes the number of images in which categories c_i and c_j co-occur. This matrix defines an undirected weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with nodes for categories and edge weights given by \mathbf{A}_{ij} . We then perform graph clustering on \mathcal{G} to divide the category set $\mathcal{C} = \{c_1, c_2, \dots, c_n\}$ into n disjoint subsets for n stages. This ensures that categories frequently co-occur are assigned to the same task, minimizing inter-task image overlap. After the above category partitioning, there still exists a portion of overlap images that contain categories from multiple stages. For each overlapping image I that spans multiple candidate category sets, we randomly assign image I to one of the candidate tasks. This strategy ensures that each image appears in only one stage, diminishing any data leakage across stages. Our proposed LoCo COCO benchmark aligns better with the real-world IOD scenario and eliminates the evaluation bias in pseudo-labeling based IOD methods.

Method

Overall Framework

The proposed YOLO-IOD utilizes the pretrained YOLO-World as its foundational model. As illustrated in Fig. 2, YOLO-IOD performs incremental learning through stage-wise partial fine-tuning, where parameter updates are guided by importance-based kernel selection. The overall architecture consists of three components. In the upper-left part, *Conflict-aware Pseudo-label Refinement (CPR)* gener-

ates high-quality old task supervision through an enhanced pseudo-label loss to mitigate foreground-background confusion from previous tasks, while addressing potential foreground-background confusion on future tasks through Clustered Unknown Pseudo Labeling. In the lower-left part of Fig. 2, *Importance-based Kernel Selection (IKS)* determines which convolutional kernels to be updated at each incremental stage, preserving critical knowledge from previous tasks while allowing partial fine-tuning to new categories. Finally, *Cross-Stage Asymmetric Knowledge Distillation (CAKD)* implements two distinct knowledge distillation losses. The first involves comparing the old detector \mathcal{M}_{t-1} with the target detector \mathcal{M}_t by feeding their respective dense features into the head of the old detector \mathcal{M}_{t-1} . The second involves comparing the current detector \mathcal{M}_{s_t} (the detector trained on current stage dataset D_t annotated for categories \mathcal{C}_t) with the target detector \mathcal{M}_t by utilizing the dense features within the current detector’s head.

Conflict-Aware Pseudo-Label Refinement

CPR comprises two components. The Enhanced Pseudo-label Loss leverages confidence-aware supervision along with entropy regularization for better use of pseudo-labels. Meanwhile, Clustered Unknown Pseudo Labeling incorporates open-vocabulary object detection and feature clustering to offer consistent supervision for objects in future tasks that are not annotated.

Enhanced Pseudo-label Loss. Traditional pseudo-labeling methods in IOD rely on confidence thresholds to select reliable pseudo labels, which introduces confidence-based selection bias. This bias causes low-confidence categories to be gradually ignored during training. In addition, pseudo labels above the threshold are treated uniformly without considering their actual reliability, resulting in inconsistent supervision signals.

To address this, we propose an Enhanced Pseudo-label Loss that treats the confidence score s of each pseudo label as a soft supervision target, integrating confidence-aware weighting with entropy regularization:

$$\mathcal{L}_{pseudo}^{cls} = -|s - p_t|^\gamma \log(p_t) + \lambda \cdot (1 - s)^\delta \cdot H(\hat{y}), \quad (1)$$

where p_t represents the predicted probability for the pseudo-labeled class, and $H(\hat{y})$ denotes the entropy of the predicted category distribution. The first term achieves confidence aligned supervision using a focal-style scheme, while the second term applies adaptive entropy regularization that scales inversely with confidence scores. This design takes full advantage of pseudo labels across wide range of confidence scores. Low-confidence pseudo labels provide soft supervision and are further regularized by the second term in Eq. (1), which retain uncertainty in the predictions, while high-confidence labels contribute more stable supervision, proportional to their reliability.

Clustered Unknown Pseudo Labeling. For the missing annotations of future task categories, we propose Clustered Unknown Pseudo Labeling method. We construct a General Vocabulary Set V_{gen} comprising 500 common object categories and 50 abstract super-categories summarized by a large language model. During each incremental stage, we

apply YOLO-World using V_{gen} to identify all foreground objects excluding those with ground-truth annotations. Let \mathcal{F} denote the set of these unannotated foreground predictions by YOLO-World, $\mathcal{C}_{\mathcal{F}} \subseteq V_{gen}$ represent the set of category labels of \mathcal{F} . To convert these predictions into stable unknown supervision while minimizing conflicts with future tasks, we perform frequency-weighted K-Means clustering on the text feature representations $f_t(\cdot)$ for $\mathcal{C}_{\mathcal{F}}$, where the weight of each category corresponds to its occurrence frequency in \mathcal{F} . The resulting cluster centroids define a set of unknown super categories $\mathcal{U} = \{u_1, u_2, \dots, u_K\}$. We then replace each label in \mathcal{F} with its assigned super-category in \mathcal{U} , and substitute their text embeddings with the corresponding cluster centroids during training. This approach transforms the knowledge conflict arising from unannotated future task categories into a process of discovering and learning new classes from the unknown super-categories.

Importance-based Kernel Selection

We utilize an importance based parameter selection mechanism to alleviate the parameter conflict in YOLO-IOD. Specifically, we select and fine-tune only important convolution kernels in each incremental task, minimizing interference with the overall parameter distribution. To avoid disrupting critical knowledge from previous tasks, we compute differential importance by subtracting historical importance from current task-specific importance.

Parameter Importance Estimation. We adopt Fisher Information to quantify parameter importance but define it at the granularity of convolution kernels to preserve inductive structure and avoid prohibitive storage costs as tasks increase. Given a convolution kernel $\mathbf{w}^k = \{w_j^k\}_{j=1}^{d_k}$ consisting of d_k scalar parameters, its Fisher-based importance can be computed as :

$$\mathbf{I}_t(\mathbf{w}^k) = \sum_{j=1}^{d_k} \left(\frac{1}{N_t} \sum_{n=1}^{N_t} \left(\frac{\partial \log p(y_n | x_n; \theta)}{\partial w_j^k} \right)^2 \right), \quad (2)$$

where (x_n, y_n) are training samples from task \mathcal{T}_t , and θ denotes the model parameters. To avoid interference with previously learned tasks, we compute the differential importance, where ρ denotes a weighting factor:

$$\Delta \mathbf{I}_t(\mathbf{w}^k) = \mathbf{I}_t(\mathbf{w}^k) - \rho \sum_{i=1}^{t-1} \mathbf{I}_i(\mathbf{w}^k). \quad (3)$$

We then select the top- \mathcal{K} kernels ranked by $\Delta \mathbf{I}_t(\mathbf{w}^k)$ for fine-tuning in task \mathcal{T}_t , while keeping the rest frozen.

Cross-Stage Asymmetric Knowledge Distillation

As illustrated in Fig. 2, our CAKD module employs a dual-teacher framework, where the target detector \mathcal{M}_t serves as the student. The first teacher is the old detector \mathcal{M}_{t-1} , which is dedicated to the previously learned categories $\mathcal{C}_{1:t-1}$ and primarily attends to the foreground features of $\mathcal{C}_{1:t-1}$, with its head suppressing responses to unrelated features. The second teacher is the current detector \mathcal{M}_{s_t} , whose features focus on the current stage classes while suppressing features

| Setting | Method | References | Baseline | AP | AP_{50} | AP_{75} | AP_S | AP_M | AP_L | $AbsGap$ | $RelGap$ | |
|-----------------|-----------------|------------|-----------------|-----------------|-------------|-------------|-------------|-------------|-------------|------------|-------------|-------|
| Upper Bound | Joint Train | NeurIPS 15 | Faster R-CNN | 40.2 | 61.6 | 43.5 | 23.5 | 43.9 | 52.2 | - | - | |
| | Joint Train | ICLR 21 | Deformable DETR | 47.0 | 66.1 | 50.9 | - | - | - | - | - | |
| | Joint Train | CVPR 24 | YOLO-World | 54.5 | 71.3 | 59.7 | 37.2 | 60.1 | 70.1 | - | - | |
| | Joint Train | CVPR 22 | GLIP-T | 55.2 | - | - | - | - | - | - | - | |
| | Joint Train | ECCV 24 | Grounding Dino | 57.2 | - | - | - | - | - | - | - | |
| 40+40 | BPF | ECCV 24 | Faster R-CNN | 34.4 | 54.3 | 37.3 | - | - | - | 5.8 | 14.4% | |
| | RGR* | CVPR 25 | Faster R-CNN | 35.6 | 56.0 | 38.2 | - | - | - | 4.6 | 11.4% | |
| | CL-DETR* | CVPR 23 | Deformable DETR | 42.0 | 60.1 | 45.9 | 24.0 | 45.3 | 55.6 | 5.0 | 10.6% | |
| | DCA | AAAI 25 | Deformable DETR | 42.8 | 58.4 | - | - | - | - | 4.2 | 8.9% | |
| | SDDGR* | CVPR 24 | Deformable DETR | 43.0 | 62.1 | 47.1 | 24.9 | 46.9 | 57.0 | 4.0 | 8.5% | |
| | TLR | AAAI 24 | GLIP-T | 40.4 | 57.4 | 43.9 | 23.3 | 44.7 | 54.5 | 14.8 | 26.8% | |
| | GCD | AAAI 25 | Grounding Dino | 45.7 | 62.9 | 49.7 | 28.4 | 49.3 | 60.0 | 11.5 | 20.1% | |
| | ERD | CVPR 22 | YOLO-Word(X) | 49.9 | 67.0 | 54.7 | 33.5 | 54.8 | 64.8 | 4.6 | 8.4% | |
| | RGR* | CVPR 25 | YOLO-Word(X) | 51.5 | 68.1 | 56.5 | 34.1 | 57.0 | 66.9 | 3.0 | 5.5% | |
| | YOLO-IOD | Ours | YOLO-Word(X) | 53.0 | 69.7 | 58.1 | 36.1 | 58.6 | 67.0 | 1.5 | 2.7% | |
| | 70+10 | BPF | ECCV 24 | Faster R-CNN | 36.2 | 56.8 | 38.9 | - | - | - | 4 | 9.9% |
| | | RGR* | CVPR 25 | Faster R-CNN | 36.6 | 56.6 | 39.6 | - | - | - | 3.6 | 8.9% |
| | | CL-DETR* | CVPR 23 | Deformable DETR | 40.4 | 58.0 | 43.9 | 23.8 | 43.6 | 53.5 | 6.6 | 14.0% |
| | | SDDGR* | CVPR 24 | Deformable DETR | 40.9 | 59.5 | 44.8 | 23.9 | 44.7 | 54.0 | 6.1 | 13.0% |
| | | DCA | AAAI 25 | Deformable DETR | 41.3 | 59.2 | - | - | - | - | 5.7 | 12.1% |
| TLR | | AAAI 24 | GLIP-T | 42.9 | 59.2 | 45.2 | 24.3 | 45.1 | 54.1 | 12.3 | 22.2% | |
| GCD | | AAAI 25 | Grounding Dino | 46.7 | 63.9 | 50.8 | 29.7 | 49.9 | 61.6 | 8.5 | 14.8% | |
| ERD | | CVPR 22 | YOLO-Word(X) | 45.6 | 62.0 | 50.1 | 29.9 | 50.6 | 60.6 | 8.9 | 16.3% | |
| RGR* | | CVPR 25 | YOLO-Word(X) | 49.1 | 64.9 | 53.5 | 31.2 | 54.1 | 64.8 | 5.4 | 9.9% | |
| YOLO-IOD | | Ours | YOLO-Word(X) | 52.4 | 68.9 | 57.4 | 35.9 | 58.0 | 65.6 | 2.1 | 3.9% | |

Table 1: Comparison with methods in single step settings. Replay-based methods are marked with *. Results are from original papers except those YOLO-World based methods which we reproduced using publicly available code. Best results in bold.

of other categories. This cross-stage asymmetric knowledge distillation design enables the target detector \mathcal{M}_t to avoid misaligned supervision and feature interference between tasks, while maximally distilling and integrating knowledge across both old and new categories.

The distillation process operates by passing student neck features $\mathbf{F}_{\text{student}}^{\text{neck}}$ to teachers’ detection heads, generating cross stage post-head features. The detection head comprises classification and regression components: the regression head outputs bounding box positions for each anchor, while the classification head produces image encodings subjected to Region-Text Matching with text embeddings to yield classification logits.

We apply distillation loss globally across the feature map. To suppress noisy or background regions and focus on the most informative and reliable predictions, we introduce a focal weight $w_{\text{focal}}(p) = \max_j \text{logit}_{\text{teacher}}(p, j)$ for each spatial location p , emphasizing foreground-likely areas based on teacher’s maximal confidence.

The classification distillation loss measures the L_2 distance between teacher and student image-region encodings at each location, weighted by the focal factor:

$$\mathcal{L}_{\text{cls_kd}} = \sum_p \|\mathbf{E}_{\text{teacher}}(p) - \mathbf{E}_{\text{student_cross}}(p)\|_2^2 \cdot w_{\text{focal}}(p), \quad (4)$$

where $\mathbf{E}_{\text{teacher}}(p)$ and $\mathbf{E}_{\text{student_cross}}(p)$ denote the region-level feature embeddings from the teacher and student models, respectively. Similarly, the regression distillation loss operates over all positions in the feature map, with the background-suppressing effect achieved via $w_{\text{focal}}(p)$:

$$\mathcal{L}_{\text{reg_kd}} = \sum_p \mathcal{L}_{\text{IoU}}(\mathbf{B}_{\text{tea}}(p), \mathbf{B}_{\text{stu_cross}}(p)) \cdot w_{\text{focal}}(p), \quad (5)$$

where $\mathbf{B}_{\text{stu_cross}}(p)$ and $\mathbf{B}_{\text{tea}}(p)$ are the bounding boxes from the student and teacher models at location p . The overall distillation objective is the weighted sum of these two parts:

$$\mathcal{L}_{\text{CAKD}} = \alpha \mathcal{L}_{\text{cls_kd}} + \beta \mathcal{L}_{\text{reg_kd}}. \quad (6)$$

Experiments

Experiment Setup

Datasets and Metrics. To evaluate our approach, we utilize the MS COCO 2017 (Lin et al. 2014) datasets with protocols established in prior studies (Mo et al. 2024; Kim et al. 2024). We also evaluate on the proposed LoCo COCO benchmark to provide more realistic assessment. We use standard COCO metrics: mAP across IoU thresholds 0.5-0.95, mAP@0.5, and mAP@0.75. We also report AbsGap

| Method | Detector | 40-10 | | | 40-20 | | | 20-20 | | | 10-10 | | |
|-------------|-----------------|-----------------|-----------------|-------------|-----------------|-----------------|-------------|-----------------|-----------------|-------------|-----------------|-----------------|-------------|
| | | \mathcal{T}_3 | \mathcal{T}_5 | RelGap | \mathcal{T}_2 | \mathcal{T}_3 | RelGap | \mathcal{T}_2 | \mathcal{T}_4 | RelGap | \mathcal{T}_4 | \mathcal{T}_8 | RelGap |
| CL-DETR* | Deformable DETR | - | 28.1 | 40.2% | - | 35.3 | 24.8% | - | 34.2 | 27.2% | - | 24.4 | 48.1% |
| SDDGR* | Deformable DETR | 40.6 | 36.8 | 21.7% | 42.5 | 41.1 | 12.5% | - | - | - | - | - | - |
| DCA | Deformable DETR | 41.1 | 37.2 | 20.8% | 42.7 | 40.3 | 14.2% | - | - | - | - | - | - |
| TLR | GLIP-T | - | 30.2 | 45.2% | - | 37.3 | 32.4% | - | - | - | - | - | - |
| GCD | Grounding Dino | - | 40.2 | 29.7% | - | 44.0 | 23.0% | - | - | - | - | - | - |
| ERD | YOLO-Word(X) | 47.3 | 37.9 | 30.4% | 50.0 | 46.4 | 14.8% | 55.1 | 44.1 | 19.1% | 46.9 | 32.0 | 41.2% |
| RGR* | YOLO-Word(X) | 48.2 | 44.8 | 17.8% | 51.0 | 48.6 | 10.8% | 56.2 | 48.1 | 11.7% | 52.3 | 43.4 | 20.3% |
| Ours | YOLO-Word(X) | 51.8 | 50.6 | 7.1% | 52.9 | 51.9 | 4.8% | 58.3 | 51.7 | 5.1% | 56.8 | 49.7 | 8.8% |

Table 2: Incremental results (AP, %) under multi-step settings. Results for 20-20, 10-10, and all YOLO-World based methods are reproduced using publicly available code. Methods using example-replay are marked with *. Best results in bold.

| Method | Baseline | 40 + 40 | | | | 70+10 | | | | 40-20 | | | |
|-------------|-----------------|-------------|------------------|------------------|-------------|-------------|------------------|------------------|-------------|-------------|------------------|------------------|-------------|
| | | AP | AP ₅₀ | AP ₇₅ | CoGap | AP | AP ₅₀ | AP ₇₅ | CoGap | AP | AP ₅₀ | AP ₇₅ | CoGap |
| RGR | Faster R-CNN | 35.0 | 55.7 | 37.2 | 0.6% | 34.6 | 54.7 | 37.4 | 2.0% | 32.5 | 52.3 | 35.0 | 1.8% |
| CL-DETR | Deformable DETR | 40.9 | 58.8 | 43.8 | 1.1% | 39.6 | 56.0 | 41.2 | 1.8% | 33.6 | 50.1 | 36.4 | 1.7% |
| GCD | Grounding Dino | 44.7 | 61.4 | 48.7 | 1.0% | 44.8 | 61.6 | 48.7 | 1.9% | 42.4 | 58.0 | 46.2 | 1.6% |
| Ours | YOLO-Word(X) | 52.2 | 68.7 | 57.3 | 0.8% | 50.7 | 67.0 | 55.6 | 1.7% | 50.9 | 66.9 | 55.7 | 1.0% |

Table 3: IOD results on LoCo COCO in single and multi-phase setting. All results are reproduced using publicly available code from their papers. *CoGap* denotes the AP gap compared to original COCO partition.

(absolute mAP Gap) and RelGap (relative mAP Gap) compared to joint training to quantify catastrophic forgetting.

Implementation Details. Our method is implemented on YOLO-World (X). We use batch size 16 on 4 RTX 3090 GPUs with learning rates of 2×10^{-5} (backbone) and 2×10^{-4} (neck and head). Training uses an AdamW (Loshchilov and Hutter 2017) optimizer for 20 epochs with mosaic augmentation disabled after epoch 10. In the IKS module, the proportion of selected kernels \mathcal{K} is set to 20% during the base stage and 12% during incremental stage.

Comparison with State-of-the-Art Methods

We assess our method under both single-step and multi-step IOD settings on COCO. We compare YOLO-IOD with recent state-of-the-art methods, including two-stage detectors (BPF, RGR), DETR-based approaches (CL-DETR, SD-DGR (Kim et al. 2024), DCA (Zhang et al. 2025)), and approaches built upon open-vocabulary models (TLR (Zhang et al. 2024), GCD (Wang, Wang, and Lin 2025)). Additionally, we reproduce classic response-based distillation method ERD and recent generative replay method RGR on the YOLO-World architecture for further comparison.

Single-Step Incremental Setting. We first evaluate performance under the 40+40 and 70+10 settings, where 40 and 10 classes are added. As shown in Tab. 1, YOLO-IOD achieves consistent performance improvements over previous methods. Specifically, under the 40+40 configuration, our method surpasses the previous best approach RGR by 1.5 in AP, with AbsGap of only 1.5, and dramatically reduces the relative performance gap to joint training from 5.5% to 2.7%. Under the 70+10 setting, YOLO-IOD out-

performs RGR by 3.3% in AP and achieves a remarkably low RelGap of 3.9% compared to the upper bound, while maintaining strong performance across all metrics. Note that RGR is a type of generative replay-based method, whereas our method requires no replay at all.

Multi-Step Incremental Setting. To better reflect real-world scenarios where new categories are continually introduced, we augment our multi-phase evaluation (40-10, 40-20) with the longer 20-20 and 10-10 settings, incrementally adding 20 or 10 classes per stage respectively until all 80 are learned. These longer configurations are especially important for evaluating incremental detectors in realistic deployment scenarios. As shown in Tab. 2, our method consistently achieves the best performance in these long-term setups, with improvement over RGR from 3.3% to 6.3% AP. Notably, under the challenging 10-10 setting with 8 incremental phases, YOLO-IOD demonstrates exceptional performance, achieving only 8.8% RelGap at the final stage and substantially outperforming RGR (20.3% RelGap) and CL-DETR (48.1% RelGap). The results indicate that YOLO-IOD can be successfully tailored for real-world scenarios demanding continuous adjustment to emerging object classes by adequately addressing knowledge conflicts.

Evaluation on LoCo COCO

We evaluate both recent methods and our YOLO-IOD on the more realistic LoCo COCO benchmark, as shown in Tab. 3. Compared to the original COCO benchmark, all methods experience an AP drop of 0.6% to 2.0% on LoCo COCO, which indicates the impact of data leakage in previous incremental settings. Despite this, YOLO-IOD consistently

| Model | Pseudo Label | CPR | IKS | CAKD | COCO(70-10) | | | COCO(40-10) | | | | | |
|-------|--------------|-----|-----|------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | | | | 1-70 | 71-80 | 1-80 | 1-40 | 41-50 | 51-60 | 61-70 | 71-80 | 1-80 |
| (a) | ✓ | | | | 49.1 | 43.4 | 48.4 | 46.5 | 33.3 | 37.9 | 56.0 | 41.2 | 44.3 |
| (b) | ✓ | ✓ | | | 51.0 | 45.4 | 50.3 | 49.9 | 36.6 | 42.7 | 58.2 | 41.2 | 47.3 |
| (c) | ✓ | ✓ | ✓ | | 52.4 | 45.2 | 51.5 | 52.8 | 36.9 | 43.6 | 57.8 | 43.6 | 49.1 |
| (d) | ✓ | | | ✓ | 51.6 | 45.3 | 50.8 | 53.0 | 36.9 | 42.6 | 58.0 | 43.9 | 49.2 |
| (e) | ✓ | ✓ | ✓ | ✓ | 53.4 | 45.4 | 52.4 | 54.5 | 38.3 | 44.6 | 58.9 | 44.6 | 50.6 |

Table 4: Investigation on the effectiveness of the main components in our method, measured in AP%.




Figure 3: Ablation study within the CAKD module on the COCO 20-20 incremental setting.




Figure 4: Ablation of kernel selection ratio \mathcal{K} in IKS on COCO 20-20 incremental setting.

achieves strong performance across all incremental scenarios, demonstrating robustness even when inter-stage image overlap is removed. Specifically, YOLO-IOD outperforms the prior best method GCD, by 7.5, 5.9, and 8.5 AP in the 40+40, 70+10, and 40+20 settings, respectively. These results highlight both the practical value of YOLO-IOD for realistic continual detection and the necessity of LoCo COCO as a practical evaluation benchmark for IOD.

Ablation Study

Effectiveness of Main components. As shown in Tab. 4, we progressively add modules to demonstrate their individual contributions in 70-10 and 40-10 settings. Starting from the standard pseudo-labeling baseline, which achieves 48.4% AP on 70-10 and 44.3% on 40-10, adding CPR significantly improves performance to 50.3% and 47.3%, respectively, demonstrating its effectiveness in mitigating foreground-background confusion. Incorporating IKS further boosts results to 51.5% AP for 70-10 and 49.1% for 40-10, high-

lighting the value of selective parameter updating. Notably, when CAKD is applied individually, AP rises from 48.4% to 50.8% in 70-10 and from 44.3% to 49.2% in 40-10, showcasing its strong distillation capability. When CAKD is combined with the previous modules, further gains can be obtained, raising AP from 51.5% to 52.4% in 70-10 and from 49.1% to 50.6% in 40-10. These results confirm that all three components CPR, IKS, and CAKD work synergistically, each contributing significant performance improvements and jointly addressing fundamental knowledge conflicts. Consequently, the complete YOLO-IOD framework achieves superior results.

Ablation Study on CAKD. As shown in Fig. 3, we compare three CAKD variants: using only the old teacher detector, only the current teacher detector, and the full dual-teacher. In early stages, the Current-Only variant performs better by facilitating fast adaptation to new categories. As tasks accumulate, the Old-Only variant becomes more effective at preserving prior knowledge, reflecting the shift from plasticity to stability. The full CAKD consistently achieves the best results, validating the advantage of asymmetric distillation that combines both knowledge sources.

Impact of IKS Kernel Selection Ratio \mathcal{K} . As depicted in Fig. 4, smaller ratios (e.g., 5%) restrict the adaptation of the model, while larger ratios (e.g., 20%) induce forgetting due to excessive updating of parameters. We observed that when setting $\mathcal{K} = 12\%$ it achieves the optimal trade-off, underscoring the significance of regulating the amount of parameters to be updated in the IOD.

Conclusions

This study introduces YOLO-IOD, a novel real-time IOD framework that leverages the pretrained YOLO-World model. By effectively addressing the challenges of foreground-background confusion, parameter interference, and misaligned knowledge distillation through the deployment of three well-designed modules: Conflict-Aware Pseudo-Label Refinement, Importance-based Kernel Selection, and Cross-Stage Asymmetric Knowledge Distillation, YOLO-IOD achieves an optimal balance between retaining prior knowledge and acquiring the ability to detect novel category objects. Furthermore, we present LoCo COCO benchmark, which successfully mitigates data leakage and aligns category partitions with real-world co-occurrences. YOLO-IOD shows outstanding performance on both conventional COCO benchmark and LoCo COCO benchmark under various single-step and multiple-step settings.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62576282, 62476223; in part by the National Key Research and Development Program of China under Grant 2024YFF1306501, 2024YFB4303700; in part by Innovation Capability Support Program of Shaanxi (Program No. 2024ZC-KJXX-043); in part by the STI2030-Major Project under Grant 2022ZD0208805; in part by the Natural Science Basic Research Program of Shaanxi Province (2024JC-DXWT-07).

References

- Asadi, N.; Davari, M.; Mudur, S.; Aljundi, R.; and Belilovsky, E. 2023. Prototype-sample relation distillation: towards replay-free continual learning. In *International conference on machine learning*, 1093–1106. PMLR.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Chen, X.; and Chang, X. 2023. Dynamic residual classifier for class incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18743–18752.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16901–16911.
- Feng, T.; Wang, M.; and Yuan, H. 2022. Overcoming catastrophic forgetting in incremental object detection via elastic response distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9427–9436.
- Kim, J.; Cho, H.; Kim, J.; Tiruneh, Y. Y.; and Baek, S. 2024. Sddgr: Stable diffusion-based deep generative replay for class incremental object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28772–28781.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; Chang, K.-W.; and Gao, J. 2022. Grounded Language-Image Pre-Training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10965–10975.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12): 2935–2947.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, 38–55. Springer.
- Liu, Y.; Schiele, B.; Vedaldi, A.; and Ruppert, C. 2023. Continual detection transformer for incremental object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23799–23808.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lu, Y.; Zhang, S.; Cheng, D.; Xing, Y.; Wang, N.; Wang, P.; and Zhang, Y. 2024. Visual prompt tuning in null space for continual learning. *Advances in neural information processing systems*, 37: 7878–7901.
- Luo, W.; Zhang, S.; Cheng, D.; Xing, Y.; Liang, G.; Wang, P.; and Zhang, Y. 2025. Gradient Decomposition and Alignment for Incremental Object Detection. In *ICCV*.
- Mo, Q.; Gao, Y.; Fu, S.; Yan, J.; Wu, A.; and Zheng, W.-S. 2024. Bridge past and future: Overcoming information asymmetry in incremental object detection. In *European Conference on Computer Vision*, 463–480. Springer.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Rypeść, G.; Cygert, S.; Khan, V.; Trzciński, T.; Zieliński, B.; and Twardowski, B. 2024. Divide and not forget: Ensemble of selectively trained experts in continual learning. *arXiv preprint arXiv:2401.10191*.
- Schlimmer, J. C.; and Granger Jr, R. H. 1986. Incremental learning from noisy data. *Machine learning*, 1(3): 317–354.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Smith, J. S.; Karlinsky, L.; Gutta, V.; Cascante-Bonilla, P.; Kim, D.; Arbelle, A.; Panda, R.; Feris, R.; and Kira, Z. 2023. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11909–11919.
- Tao, X.; Chang, X.; Hong, X.; Wei, X.; and Gong, Y. 2020. Topology-preserving class-incremental learning. In *European conference on computer vision*, 254–270. Springer.

Von Oswald, J.; Henning, C.; Grewe, B. F.; and Sacramento, J. 2019. Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695*.

Wang, X.; Wang, Z.; and Lin, Z. 2025. GCD: Advancing Vision-Language Models for Incremental Object Detection via Global Alignment and Correspondence Distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 8015–8023.

Wang, Z.; Zhang, Z.; Lee, C.-Y.; Zhang, H.; Sun, R.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 139–149.

Wu, Q.; Zhang, S.; Cheng, D.; Xing, Y.; Xu, D.; WANG, P.; and Zhang, Y. 2025. Demystifying Catastrophic Forgetting in Two-Stage Incremental Object Detector. In *Forty-second International Conference on Machine Learning*.

Yang, Y.; Cui, Z.; Xu, J.; Zhong, C.; Zheng, W.-S.; and Wang, R. 2023. Continual learning with bayesian model based on a fixed pre-trained feature extractor. *Visual Intelligence*, 1(1): 5.

Zhang, A.; Yang, D.; Liu, C.; Hong, X.; Shang, M.; and Zhou, Y. 2025. DCA: Dividing and Conquering Amnesia in Incremental Object Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9851–9859.

Zhang, H.; Gao, B.-B.; Zeng, Y.; Tian, X.; Tan, X.; Zhang, Z.; Qu, Y.; Liu, J.; and Xie, Y. 2024. Learning task-aware language-image representation for class-incremental object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 7096–7104.