

DitHub: A Modular Framework for Incremental Open-Vocabulary Object Detection

Chiara Cappellino Gianluca Mancusi Matteo Mosconi Angelo Porrello
Simone Calderara Rita Cucchiara

AImageLab - University of Modena and Reggio Emilia

name.surname@unimore.it

Abstract

Open-Vocabulary object detectors inherently generalize to an unrestricted set of categories, enabling recognition through simple textual prompting. However, adapting these models to rare classes or reinforcing their abilities on specialized domains remains essential. While recent methods rely on monolithic adaptation strategies with a single set of weights, we embrace modular deep learning. We introduce DitHub, a framework designed to build and maintain a library of efficient adaptation modules. Inspired by Version Control Systems, DitHub manages expert modules as branches that can be fetched and merged as needed. This modular approach allows us to conduct an in-depth exploration of the compositional properties of adaptation modules, marking the first such study in Object Detection. Our method achieves state-of-the-art performance on the ODinW-13 benchmark and ODinW-O, a newly introduced benchmark designed to assess class re-appearance. For more details, visit our project page: <https://aimagelab.github.io/DitHub/>

1. Introduction

Object detection is a cornerstone of computer vision, with recent advancements [1, 20, 26, 42, 51] enabling architectures that integrate Vision and Language large-scale pre-training [21]. This has given rise to *Open-Vocabulary Object Detection*, where models can classify and localize objects beyond a predefined set of categories. The ability to process arbitrary textual queries and detect corresponding objects has profound implications for real-world applications, including medical imaging [15], safety and security [35], and industrial quality control [32]. For instance, the ability to detect rare or atypical objects from a safety camera is crucial for issuing early warnings about uncommon dangers. However, this requires Open-Vocabulary de-




Figure 1. DitHub: a novel system to share and update deep adaptation modules for Open-Vocabulary Object Detection.¹

tectors to adapt effectively, enabling them to recognize categories beyond their pre-training knowledge distribution.

Recent approaches, such as those detailed in [2], have extended vision-language detectors to **incrementally** accommodate new categories while preserving robust zero-shot capabilities [34]. However, these methods often employ *monolithic* adaptation, where all newly acquired knowledge is condensed into a single set of weights. This approach presents challenges in real-world scenarios that demand updates to specific concepts which may reappear under varied input modalities. For instance, in safety and security applications, the class “person” may need to be detected not only in standard RGB imagery but also in thermal imagery, necessitating continuous model adaptation to ensure consistent performance across different data types. Additionally, for rare or complex categories, incremental adaptation

¹The creature in the red box is an axolotl, symbolizing rarity as a deviation from pre-training. We use a stylized axolotl as the icon for DitHub.

is crucial to refine the base model for fine-grained concepts. In such scenarios, a monolithic deep architecture faces challenges similar to those of managing a complex program written on a single page of code. Namely, when treating multiple classes within a unique set of weights, it becomes hard to update specific concepts selectively, without compromising information related to other concepts. Considering less-represented categories, the corresponding concepts can become blurred within the unique set of weights, leading to potential degradation in performance over time.

We advocate for **modular** approaches, as outlined in Modular Deep Learning [30], which reimagine neural networks as adaptable and composable entities rather than rigid, monolithic structures. By maintaining a library of expert modules, neural architectures can achieve class-specific specialization, activating only the relevant module when needed. Moreover, this modular design enables selective updates and seamless composability, allowing class-specific modules trained in different domains to be efficiently merged. Finally, modularity offers a practical solution to mitigate biases — if a module exhibits undesirable biases toward a particular category, it can be selectively removed or adjusted without disrupting the entire model, ensuring more precise and fair adaptation.

With modularity at its core, we introduce DitHub (where D stands for Detection), a framework inspired by Version Control Systems that efficiently manages a growing library of detection modules. This library expands incrementally as new groups of categories, referred to as tasks [41], are introduced. Upon encountering a new task, DitHub initializes an efficient module (e.g., LoRA [11]), which first undergoes a broad pre-tuning phase across all classes of the task, referred to as the **warmup** phase. This stage establishes a solid and class-agnostic foundation, enabling the subsequent **specialization** of the module for specific classes. Following this phase, the module *branches* into specialized experts, each optimized independently. To enable selective updates for classes previously introduced to the model, DitHub first performs a *fetch* operation to access potentially existing class-specific modules. If these exist, instead of starting the training from the warmup, DitHub uses a *merged* initialization — combining the retrieved pre-existing class-specific modules with the current warmup module. This merged initialization facilitates smooth domain knowledge transfer and ensures continuity of class-specific learning across tasks.

Moreover, to address potential memory constraints, DitHub employs a subset of shared learnable parameters across different adaptation modules, enhancing efficiency without compromising individual class adaptability. At inference, it enables the selective activation of specialized modules for rare or complex classes, ensuring robust adaptation while maintaining the backbone’s zero-shot and

open-vocabulary capabilities. A high-level overview of our approach is depicted in Fig. 1.

In the context of [2], we evaluate DitHub on the ODinW-13 [17] benchmark, achieving state-of-the-art performance within the Incremental Vision-Language Object Detection paradigm. To further assess our proposal, we introduce ODinW-O (Overlapped), a curated subset of ODinW-35 [16] that focus on classes re-appearing across multiple tasks. Our simple yet effective methodology — without additional loss terms or architectural modifications — allows us to investigate the compositional capabilities of efficient modules, marking the first exploration of this kind in Object Detection. In summary, our key contributions are:

- We introduce DitHub, a framework inspired by Version Control Systems, designed to manage a growing library of efficient modules for Open-Vocabulary detection.
- We achieve state-of-the-art performance on ODinW-13 and the newly introduced ODinW-O.
- We conduct an in-depth analysis of the compositional and specialization capabilities of class-specific modules, a first in the context of Object Detection.

2. Related Works

Incremental Vision-Language Object Detection

Vision-Language Object Detection (VLOD), also referred to as Open-Vocabulary detection [2], is the task of detecting and recognizing objects beyond a fixed set of predefined categories. Early approaches leverage powerful visual and textual encoders (e.g., CLIP [34]) to enable object detectors to recognize virtually any object specified by a textual query [8, 49]. A fundamental trend in this field has been the reliance on large-scale pre-training on extensive datasets [37], which has become standard practice for tasks such as classification. In the context of Vision-Language Object Detection, state-of-the-art methods have followed this paradigm, with GLIP [17] being one of the first notable examples. GLIP reframes object detection as a phrase-grounding problem, leading to a pre-trained model capable of generalizing to unseen objects by integrating textual and visual semantics from large-scale datasets. Building on GLIP’s problem formulation, Grounding DINO [21] extends this paradigm by incorporating merged visual and textual semantics at multiple network stages.

Recently, Incremental Vision-Language Object Detection (IVLOD) [2] has emerged as a natural extension of the VLOD paradigm. By incorporating Incremental Learning [14, 36, 41, 45], IVLOD addresses the challenge of fine-tuning a pre-trained Vision-Language model on specific object categories while preserving its zero-shot [29] capabilities. Following the approach of [2], we adopt Grounding DINO as our backbone and utilize IVLOD as the foundation to explore the merging of class-specific modules.

Parameter Efficient Fine-Tuning and Model Merging

The rise of Transformer models [5, 42] has made pre-training a standard approach in machine learning. Parameter-Efficient Fine-Tuning (PEFT) methods, such as Prefix Tuning [18], Adapters [10], and LoRA [11], adapt pre-trained models to downstream tasks by updating only a small fraction of parameters while keeping the rest frozen.

In the context of Incremental Vision-Language Object Detection, ZiRa [2] introduces a distinct approach that can be viewed as a specialized form of parallel adapters [23], incorporating reparameterizable side branches. In contrast, our work adopts LoRA and explores the potential of merging class-specific LoRA modules.

Several studies have investigated LoRA module fusion. An early approach relies on simple arithmetic operations such as addition and subtraction [50]. More advanced techniques focus on learning optimal combination coefficients, determining the contribution of each LoRA module in the merging process [12, 46–48]. Recently, researchers have introduced closed-form solutions for LoRA merging [38] and proposed regularization objectives to ensure that the fused modules remain close to the pre-trained weight distribution [31]. In this respect, our approach takes a novel direction, being the first to study module merging in the context of object detection.

3. Preliminaries

Problem setting. We consider a sequence of tasks $\{1, \dots, T\}$, where each task t corresponds to a dataset D_t . Each dataset consists of samples (x, y) , where x represents the input data, and y contains the associated labels, including ground truth bounding boxes.

Given a pre-trained Vision-Language detector $f(\cdot; \Theta)$, parameterized by Θ , the objective of *Incremental Vision-Language Object Detection* (IVLOD) is to improve performance on newly introduced tasks while preserving original zero-shot capabilities. Formally, for a given task t , we seek an adaptation Δ that minimizes the following objective:

$$\min_{\Delta} \mathbb{E}_{(x,y) \sim D_t} \ell(f(x; (\Theta, \Delta)), y) + \ell_{\text{zero-shot}}, \quad (1)$$

where ℓ is the object detection loss function. The term (Θ, Δ) represents the set of adapted model parameters. Additionally, $\ell_{\text{zero-shot}}$ quantifies the performance difference in the zero-shot setting before and after adaptation.

Each task t introduces a set of classes $\mathcal{C}_t$. Unlike standard incremental learning [40, 41, 44, 45], IVLOD allows class overlap between tasks. In other words, there exist task pairs (t, t') , with $t' \neq t$, such that $\mathcal{C}_t \cap \mathcal{C}_{t'} \neq \emptyset$.

LoRA. Low-Rank Adaptation [11] is a technique to significantly reduce the number of trainable parameters. Given a generic layer with weight matrix $W \in \mathbb{R}^{d \times k}$, LoRA introduces low-rank matrices $A \in \mathbb{R}^{r \times k}$ and $B \in \mathbb{R}^{d \times r}$ — with

$r \ll \min(d, k)$ — such that the adapted weight W' is:

$$W' = W + \Delta W, \quad \text{where} \quad \Delta W = BA. \quad (2)$$

The use of these low-rank matrices significantly reduces memory usage and computational overhead, making LoRA an efficient alternative to fine-tuning the entire model.

4. DitHub

Version Control Systems are ubiquitous tools in computer science, renowned for their capacity to efficiently manage code repositories. Inspired by this capability, we introduce DitHub — a framework that operates akin to a version control system, facilitating the creation and continuous expansion of a library of efficient adaptation modules for vision-language object detectors. To develop and maintain class-specific *expert* modules, DitHub employs a straightforward yet effective strategy that requires no architectural modifications or additional loss terms, as detailed in Sec. 4.1.

In line with the authors of [2], we consider the pre-trained Grounding DINO as the base architecture for Open-Vocabulary detection. We employ LoRA as a fine-tuning strategy, proposing a novel and efficient pipeline that handles its low-rank matrices with two specific objectives. In our approach, matrix A is dedicated to class-specificity, while matrix B captures general knowledge and is shared among tasks to maintain overall *efficiency*, as further discussed in Sec. 4.2. Finally, when a specialized task arises or more fine-grained adaptation is required (e.g., detecting a rare object within a scene), DitHub enables the dynamic extraction of the corresponding expert modules, ensuring precise and efficient inference, as outlined in Sec. 4.3. The pseudocode of DitHub is provided in Alg. 1.

4.1. Training Class-Specific Modular Detectors

To manage a library of class-specific modules, the approach is to train a tailored A matrix for each class, allowing it to specialize accordingly. On the contrary, the matrix B is shared across classes and continually fine-tuned while tasks progress. Thus, at task t , we train a total of $|\mathcal{C}_t| + 1$ matrices for each layer that requires adaptation. The total includes a matrix for each of the $|\mathcal{C}_t|$ classes introduced in the current task, along with a single matrix for B .

However, in object detection, creating class-specific experts is non-trivial, as a single image may contain multiple bounding boxes spanning different categories, making it difficult to focus on a certain class. To tackle this, we adopt a **stochastic training strategy** for specializing the A matrices: given an image with multiple objects from different classes, we randomly select one of the present classes and execute a training step on its corresponding class-expert module. However, in our preliminary experiments, this strategy proved ineffective for rare categories, which may




Figure 2. Overview of DitHub. Through three primitives — *i.e.*, branch, fetch, merge — DitHub allows to update a library of class-specific modules for Open-Vocabulary Object Detection, ensuring a balanced integration of both past and present knowledge.

appear only a few times in real-world scenarios. Relying solely on random selection prevents the A matrix of rare classes from accessing enough training data.

As better highlighted in Alg. 1, we overcome this issue by introducing an earlier **warmup** phase at the beginning of each task. This phase is meant to provide a robust, class-agnostic initialization, ensuring that the subsequent **specialization** stage is more effective.

Warmup. At the onset of the new task, a single and shared warmup matrix A_{wu} is instantiated and optimized (see Fig. 2, yellow arrow) to capture class-agnostic yet task-specific knowledge. This phase, spanning a predefined number of epochs, establishes a common initialization for the A modules, which are set to undergo specialization in the subsequent stage. By doing so, the specialized modules start from a common base point, a factor that has been shown to greatly enhance model compositionality [27, 33]. Often associated with improved linear mode connectivity between models [7], these properties are leveraged by our approach to select and combine class-specific modules based on input textual queries, as detailed in Sec. 4.3. Moreover, considering rare and underrepresented classes with a few training examples, the warmup phase seeks to establish an already robust and effective initialization.

Specialization. Once the warmup phase is finished, we begin with the specialization phase, the cornerstone of DitHub. To train class-specific modules, DitHub generates $|\mathcal{C}_t|$ distinct branches (branch in Fig. 2), each dedicated to a single class. These branches undergo independent optimization, represented by a blue arrow in Fig. 2, to capture class-specific capabilities. In doing so, for each training image, we employ the stochastic strategy outlined at the beginning of Sec. 4.1 to randomly select a single expert for training.

Notably, the modular design of our approach effectively handles potential *class overlaps* across consecutive tasks. In real-world scenarios, classes may reappear in new domains with altered distributions or semantics. For instance, consider the class “dog”, originally learned from real-world camera images, later appearing in a task involving thermal imagery. The goal is not only to retain prior knowledge but also to develop a “dog” expert capable of recognizing the entity in both domains, accommodating any semantic shifts. In this case, the modularity of DitHub enables the selective adaptation of the “dog” module, while minimizing interference with other classes, such as the “axolotl” module.

To favor selective updates, DitHub performs a *fetch* operation (Fig. 2) to retrieve any previously stored version of the specialized module, such as A_{dog}^{old} . If found, this module is merged (merge in Fig. 2) with the current task’s warmup matrix A_{wu} . For instance, considering the class “dog”, the merge operation follows a weighted strategy:

$$A_{dog}^{cur} = (1 - \lambda_A)A_{dog}^{old} + \lambda_A A_{wu} \quad (3)$$

Here, λ_A is set to a low value, prioritizing consolidated knowledge over new insights from the warmup phase. Details on the impact of λ_A are in Sec. 5.3. Crucially, our fetching process creates an on-the-fly initialization that blends historical class-specific knowledge with new domain information.

4.2. Bridging Modularity with Memory Efficiency

To keep our library memory-efficient, we maintain a single shared B matrix instead of a distinct B for each class. As shown in Fig. 3, using class-specific (A, B) pairs results in a linear increase in memory consumption as the number of classes grows. Sharing the B matrix reduces

Algorithm 1 DitHub Training

Input: T disjoint tasks each containing $|\mathcal{C}_t|$ classes; library $\mathcal{A} \leftarrow \{\}$ of class experts; global B .

for task $t \in \{1, \dots, T\}$ **do**

$(A_{\text{wu}}, B^{\text{opt}}) \leftarrow \text{warmup (training)}$

for class $c \in \mathcal{C}_t$ **do** ▷ Branch

if $A_c \in \mathcal{A}$ **then** ▷ Fetch

$A_c^{\text{cur}} \leftarrow \text{use Equation 3}$ ▷ Merge

else

$A_c^{\text{cur}} \leftarrow A_{\text{wu}}$

$(A_c, B^{\text{opt}}) \leftarrow \text{specialization (training)}$

$\mathcal{A} \leftarrow \mathcal{A} \cup A_c$

$B \leftarrow \text{use Equation 4}$




Figure 3. Memory requirements as the number of classes grows.

this requirement by half². Indeed, considering DitHub “ B shared” (line $\circ$), the memory footprint of our approach is consistently lower than the storage needed for the Grounding DINO encoder model *alone* (line $---$). This makes DitHub’s memory demands manageable for practical applications. Also, the memory efficiency is further optimized by fine-tuning only a limited number of layers (see App. A).

Due to the shared design, the matrix B is more susceptible to interference and catastrophic forgetting [25]. Consequently, we implement a tailored mechanism that ensures smooth updates as tasks progress. Specifically, at the end of each task t , we update the matrix B by merging the previous checkpoint, B_{t-1} , with the weights learned after the current task, B^{opt} . This process is formulated as follows:

$$B_t = (1 - \lambda_B)B_{t-1} + \lambda_B B^{\text{opt}}. \quad (4)$$

Unlike Eq. 3, the merging coefficient λ_B is assigned a higher value, giving greater weight to the latest state. This approach results in a dual strategy for λ_A and λ_B , which will be further explained in Sec. 5.3.

²In App. A.2, we present DitHub’s results when B is not shared.

4.3. Detection with Class-Specific Modules

With our library of efficient, class-specific modules in place for adapting the Vision-Language backbone, two disjoint inference routes are available for users. Given a textual query specifying the classes of interest, our model tokenizes the prompt and searches token by token to determine if corresponding expert modules are present in the library. If there are correspondences, those relevant modules are fused using a simple average to create a single composite adaptation module. Conversely, for requests targeting individual classes where fine-grained adaptation is desired, the corresponding module can be directly selected and deployed. In our experimental studies, we present results and discussions that address both scenarios.

5. Experimental Studies

This section evaluates DitHub in two parts. The first assesses its adaptability to multiple tasks using the Incremental Vision-Language Object Detection (IVLOD) protocol [2] in Sec. 5.1. The second, in Sec. 5.2, examines the method’s specialization capabilities. The section ends with ablation studies in Sec. 5.3.

5.1. Incremental Vision-Language Object Detection

Evaluation settings. Following [2], we employ ODinW-13 [17] to assess the model ability to learn new tasks while mitigating forgetting of previously seen data. ODinW-13 consists of 13 sub-datasets, spanning from the traditional Pascal VOC [6] to more challenging domains with significant distribution shifts, such as thermal imagery. Also, the IVLOD setting mandates that the Open-Vocabulary detector maintains robust zero-shot capabilities after fine-tuning. These capabilities are tested on MS COCO [19].

One issue of ODinW-13 is that most classes are confined to a single task among the 13. However, DitHub is conceived to handle classes that reoccur across different domains, necessitating updates to their respective modules. To provide an additional perspective on incremental object detection and model reuse, we introduce a variant of ODinW-13, referred to as **ODinW-O**. Starting from ODinW-35, we construct ODinW-O by retaining only the classes that are shared across tasks while discarding those that appear in only one task. Further details are provided in App. B.

For comparisons, we employ mean Average Precision (mAP) and compare DitHub against six state-of-the-art methods. CL-DETR [22] and OW-DETR [9] are explicitly designed for incremental learning, while TFA [43] specializes in few-shot adaptation. AT [10] has been extended to both incremental and few-shot learning, whereas iDETR [4] integrates both aspects. Finally, ZiRa [2], our primary competitor, is specifically tailored for IVLOD.

Implementation details. We focus on Grounding DINO

Shots	Method	ZCOCO	Avg	Ae	Aq	Co	Eg	Mu	Pa	Pv	Pi	Po	Ra	Sh	Th	Ve
0	G-Dino	47.41	46.80	19.11	20.82	64.75	59.98	25.34	56.27	54.80	65.94	22.13	62.02	32.85	70.38	57.07
1	TFA	18.84	39.50	18.25	15.81	63.90	50.79	28.47	50.37	29.49	59.16	21.90	50.67	19.86	60.85	43.97
	iDETR	44.61	49.82	22.81	23.24	69.75	61.43	31.73	56.27	55.40	62.44	28.45	60.33	43.33	73.64	58.84
	AT	44.11	46.23	21.55	23.62	66.60	58.96	27.68	53.97	54.58	62.47	26.94	53.17	20.37	70.31	60.71
	ZiRa	46.44	48.56	20.34	19.64	69.47	60.00	30.09	56.27	58.12	65.13	25.83	55.69	39.50	72.59	58.55
	DitHub	46.40	49.19	21.33	24.34	69.23	61.99	42.10	57.08	58.85	55.24	24.62	57.58	33.56	70.99	62.50
5	TFA	26.43	45.76	21.92	22.30	67.40	60.72	30.63	53.56	46.80	63.60	26.88	56.26	28.00	64.28	52.49
	iDETR	43.51	51.65	25.69	25.53	70.42	62.98	49.98	50.54	54.85	64.80	33.24	57.64	42.36	76.51	56.92
	AT	43.67	47.16	14.63	24.97	66.56	64.19	38.85	42.03	55.49	65.20	27.48	52.68	32.21	71.23	57.54
	ZiRa	45.41	51.77	24.22	26.40	71.34	63.96	49.54	60.51	60.11	64.39	32.20	56.31	40.61	69.16	54.30
	DitHub	45.86	52.85	26.84	28.91	68.46	60.25	51.93	55.10	60.43	59.44	33.18	68.12	38.05	75.04	61.34
10	TFA	34.35	46.61	21.17	22.16	66.82	60.63	32.35	50.15	55.54	64.98	27.59	57.40	28.14	66.82	52.11
	iDETR	43.54	53.29	25.39	27.70	65.62	67.58	47.99	60.20	56.32	63.93	35.18	59.37	53.63	74.70	55.12
	AT	43.06	47.34	18.73	25.42	69.77	66.34	35.84	48.25	53.39	64.07	28.89	50.33	32.50	66.57	55.37
	ZiRa	46.14	53.20	25.56	25.96	71.72	65.28	48.22	62.81	61.24	65.47	35.58	55.73	46.86	70.97	56.14
	DitHub	46.54	54.43	26.19	30.28	68.11	65.73	52.82	56.82	63.36	61.85	36.68	66.82	41.46	76.00	61.48
Full	TFA	30.97	47.93	23.80	30.65	67.21	61.77	30.52	50.23	47.73	60.91	29.25	61.72	31.42	66.23	61.61
	iDETR	37.32	58.71	32.64	46.65	70.99	68.56	55.32	58.88	64.48	71.01	50.33	63.30	39.19	77.12	64.80
	AT	42.30	51.14	23.62	39.90	72.32	65.51	31.47	50.48	60.51	66.07	39.09	53.50	34.04	68.07	60.23
	OW-DETR	31.22	55.58	28.46	43.78	70.54	67.78	43.84	56.75	63.13	69.51	45.16	58.99	36.99	74.42	63.20
	CL-DETR	32.15	57.26	29.35	45.15	71.94	69.90	45.21	58.52	65.10	71.68	46.58	60.83	38.14	76.74	65.18
	ZiRa	46.26	57.98	31.76	47.35	71.77	64.74	46.53	62.66	66.39	71.00	48.48	63.03	41.44	76.13	62.44
	DitHub	47.01	62.19	34.62	50.65	70.46	68.56	49.28	65.57	69.58	71.10	56.65	70.88	52.82	79.30	68.18

Table 1. Comparison of mAP values across ODinW-13 tasks. **Avg** represents the average across the 13 tasks, while zero-shot performance is reported in the **ZCOCO** column. Best results are highlighted in bold.

and adapt only its encoder network, as further validated in App. A.1. The merging coefficient λ for the A matrix is set to 0.3 for ODinW-13 and 0.1 for ODinW-O; for the matrix B , we set $\lambda = 0.7$. The choice of λ and its underlying rationale are further discussed in Sec. 5.3. During training, we allocate an equal number of epochs to the warmup and the subsequent specialization phases — App. C for further hyperparameter details.

Performance on ODinW-13. Table 1 presents the IVLOD results, including the average across the 13 tasks (Avg) and zero-shot preservation (ZCOCO). We evaluate performance in both the standard setting with the full dataset and the few-shot scenarios, considering 1, 5, and 10 shots per class. All baseline results, except those of Grounding DINO, ZiRa, and our method, are taken directly from [2]. Each number represents the average of three runs with different random seeds, with standard deviations provided in App. D.1.

In the full-shot setting, DitHub outperforms its main competitor, ZiRa, by a substantial + 4.21 mAP points and achieves the best results in nine out of thirteen tasks. Notably, on ZCOCO, our approach outperforms ZiRa by a mAP gain of + 0.75, setting a new state-of-the-art in both incremental and zero-shot retention capabilities. Although DitHub is not explicitly designed for data-scarce conditions, DitHub achieves state-of-the-art performance even in the few-shot setting. In particular, in the five-shot and ten-shot settings, DitHub emerges as the best-performing approach, with the ten-shot scenario showing a slightly larger

	ZCOCO	Avg	Ae	Hw	Pv	Sd	Th	Ve
G-Dino	47.41	53.15	45.12	67.54	58.11	25.84	70.40	51.87
ZiRa	44.43	57.63	39.92	68.00	64.90	46.26	77.26	49.47
DitHub	46.51	62.38	53.35	71.07	71.01	41.75	80.21	56.90

Table 2. mAP values on ODinW-O. Best in bold.

improvement over ZiRa (+ 1.23 mAP). Other baselines struggle to remain competitive in this regime. Our method also performs well in the one-shot setting, where iDETR achieves the highest average performance, followed closely by DitHub and ZiRa.

These findings — optimal performance in incremental learning, effective retention of zero-shot capabilities, and robustness in the low-data regime — clearly highlight the advantages of our modular approach, with the advanced update system that will be dissected in the analyses below.

Performance on ODinW-O. The ODinW-O scenario emphasizes that frequently encountered classes may reappear in subsequent tasks or at later times, such as the class “person” in common industrial and smart-city scenarios. As detailed in App. B.2, ODinW-O guarantees that each class is revisited at least once in subsequent tasks. In this benchmark, we primarily compare our approach with ZiRa, as it is the most competitive method in the standard IVLOD setting, excluding ours. The results in Tab. 2 indicate that our approach secures a substantial improvement of + 4.75 in

Method	ZCOCO	Avg	Ae	Aq	Co	Eg	Mu	Pa	Pv	Pi	Po	Ra	Sh	Th	Ve
G-Dino	47.41	46.80	19.11	20.82	64.75	59.98	25.34	56.27	54.80	65.94	22.13	62.02	32.85	70.38	57.07
EnE	46.86	60.96	40.51	44.88	69.67	67.54	51.62	63.42	70.32	70.29	53.76	69.93	45.19	78.36	66.95
DitHub	47.01	62.19	34.62	50.65	70.46	68.56	49.28	65.57	69.58	71.10	56.65	70.88	52.82	79.30	68.18

Table 3. Experts not Experts (EnE). Training by selecting a random expert.




Figure 4. Specialization capabilities for rare classes of ODiW-13.

mAP and a **+ 2.08** mAP gain on ZCOCO over ZiRa. We attribute the success of our method to its class-oriented modular design, which enables selective updates to recurring concepts. This design effectively prevents the model from inadvertently overwriting crucial knowledge associated with other classes, tasks, or domains.

5.2. Evaluation of Specialized Detection Modules

Here, we delve into the properties of our specialization framework and the benefits of composing class-specific experts. To this aim, we introduce an ablative approach called **Experts-not-Experts (EnE)**, which emphasizes the absence of class specialization in our experts. EnE is identical to DitHub, except during training, where it does not employ the stochastic strategy introduced in Sec. 4.1 and designed to specialize the experts. Instead, given an image from the current task, the EnE approach performs the training step on one module chosen at random among the classes of the entire task. As a result, modules no longer specialize in a class but instead learn from the whole task data, including images without their designated class. Consequently, they become general task learners rather than class-specific experts.

In Tab. 3, we offer a comparative analysis between our DitHub and its non-specialized variant, *i.e.*, EnE. Notice that, at inference time, EnE acts as an ensemble of modules, each trained on the same data but with a different random ordering. In contrast, DitHub employs class-specialized modules with an expertise on data containing the same class. This targeted specialization improves performance, as shown in Tab. 3, where per-class specialization results in a **+ 1.23** mAP improvement on average. Additionally, albeit to a lesser extent, the ZCOCO metric still benefits from an improvement. These results affirm that a system composed




Figure 5. Showing the natural unlearning capabilities of DitHub.

of modular experts, each a custodian of distinct knowledge, offers a promising approach for systems requiring continuous and selective adaptation.

Evaluation on rare classes. As discussed in Section 4.3, DitHub facilitates fine-grained control during inference, enabling users to tailor the system to their specific use case and deploy the most appropriate modules. This includes the option to use a single module specialized on a target class. To assess single-module effectiveness, we conduct a targeted evaluation on rare classes from ODiW-13, activating only the corresponding target class module. Results in Fig. 4 show that DitHub consistently achieves the best results across rare classes, highlighting the benefits of class specialization. EnE and DitHub outperform the zero-shot Grounding DINO baseline, though EnE underperforms for the Starfish category.

Specialized modules allows training-free unlearning. The advantages of DitHub extend beyond additive incremental knowledge compositionality. Indeed, we can show that expert modules can be useful to undermine or remove some specific knowledge or content from a model, by simply subtracting the corresponding weights from the base model. This operation, which resembles the concept of *zero-training unlearning*, has been explored by the authors of [13, 28] in the context of image and text categorization tasks; to the best of our knowledge, we are the first to explore its application within the context of object detection.

Specifically, we test whether class knowledge can be unlearned from the zero-shot Grounding DINO baseline by




Figure 6. Sensitivity analysis on λ_A and λ_B . Performance measured as harmonic mean of Avg mAP (across tasks) and ZCOCO.

subtracting the corresponding class module from the pre-trained model: $\mathbf{W}' = \mathbf{W} - \alpha \mathbf{B} \mathbf{A}_c$, where we set $\alpha = 0.3$. As a result, we provide a qualitative visualization shown in Fig. 5: on the left, the zero-shot model correctly recognizes common classes, while on the right, applying the negative module of a specific class reduces its mAP, effectively unlearning it and leading to false negatives. Additionally, we provide a quantitative evaluation of the unlearning process in App. D.2.

5.3. Ablation Studies

Sensitivity Analysis of λ . Analyzing the variation of the merging coefficients λ_A and λ_B offers valuable insights into our methodology. To show that, we define an aggregate metric that is the harmonic mean between DitHub’s full-shot performance on ZCOCO and Avg, ensuring a balanced evaluation of specialization (Avg) and zero-shot generalization (ZCOCO). Figure 6 presents the performance trends for ODinW-13 (top) and ODinW-O (bottom) while varying λ_A and λ_B in increments of 0.2.

Considering λ_A , its role is significant while merging the expert A matrices (see Eq. 3). From Fig. 6, lower values of λ_A lead to better results, indicating that prioritizing past expert knowledge over the warmup module is more effective. Interestingly, while $\lambda_A = 0.3$ performs best for ODinW-13, ODinW-O exhibits a monotonic decline as λ_A increases. This contrast arises from dataset structures: ODinW-O consists solely of overlapping classes across tasks, where retaining past expertise is critical, whereas ODinW-13 features diverse class distributions per task, requiring greater integration of domain-specific warmup knowledge.

For the incremental adaptation of the shared B matrix, a different trend emerges. The harmonic mean increases with λ_B , peaking at 0.7 for both datasets. This suggests that prioritizing the knowledge in the B^{opt} matrix from the current task, as shown in Eq. 4, is crucial. Our intuition, supported by [24], is that the latest B^{opt} matrix integrates knowledge from both earlier and recent tasks, making it more valuable




Figure 7. Ablative study on the DitHub components. Results are expressed as average performance across the tasks of ODinW-13.

than B_{t-1} , which is trained solely on earlier tasks.

On Warmup and Shared B . To evaluate the impact of each component in DitHub, we progressively integrate them and present the results in Fig. 7. We start with the zero-shot performance of Grounding DINO (GD) as a reference, followed by a baseline fine-tuning procedure (Base in Fig. 7), where an A matrix is trained for each class without merging strategies, omitting the warmup phase and using a shared B matrix. This method leads to complete forgetting as new tasks are introduced, but still outperforms Grounding DINO, which struggles with pretraining knowledge alone.

Building on this baseline, we first incorporate Eq. 4, but see no significant performance gain. This is expected, as the shared B matrix only represents a tiny portion of the total parameters. Without class-specific focus, Eq. 4 alone does not enhance retention. However, adding the warmup phase (+wu) provides a clear improvement, as it stabilizes training and helps rare classes by offering a robust initialization.

When we further include Eq. 4 in the warmup-enhanced fine-tuning strategy, a modest improvement is observed. However, the most substantial performance gain occurs when Eq. 3 is also integrated into the methodology (DitHub in Fig. 7). This highlights the importance of preserving past domain knowledge within class-specific modules, ensuring a more stable and effective incremental learning process. Notably, the combination of all three key components — the warmup phase for reliable initialization, the merging strategy for A matrices to retain class-specific knowledge across domains, and the shared B matrix approach for efficiency — results in a significant + 6.22 improvement over the baseline where none of these strategies are applied.

6. Conclusions

In this research, we introduced DitHub, a methodology akin to a version control system that maintains an expanding library of modules to adapt open-vocabulary object detectors. Our approach achieves state-of-the-art performance in the Incremental Vision-Language Object Detection scenario. The incremental nature of this domain allowed us to explore the compositional capabilities of efficient modules, marking the first such investigation in Object Detection.

Looking ahead, we aim to further leverage DITHub’s modular library management to deepen our understanding of knowledge transfer between efficient modules for object detection. For instance, if our library contains an “axolotl” module trained on real-world images and we possess domain knowledge for the cartoon image domain, can we synthesize a module specialized in detecting cartoon “axolotls” through knowledge transfer? DITHub will provide a fork of this manuscript to explore such questions.

Acknowledgments

This article is the result of a collaboration with the partners of the STORE project (<https://edf-store.com>), funded by the European Defence Fund (EDF) under grant agreement EDF-2022-101121405-STORE.

References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision*, 2020. 1
- [2] Jieren Deng, Haojian Zhang, Kun Ding, Jianhua Hu, Xingxuan Zhang, and Yunkuan Wang. Zero-shot generalizable incremental learning for vision-language object detection. In *Advances in Neural Information Processing Systems*, 2024. 1, 2, 3, 5, 6, 12
- [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2019. 12
- [4] Na Dong, Yongqiang Zhang, Mingli Ding, and Gim Hee Lee. Incremental-detr: Incremental few-shot object detection via self-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 5
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 3
- [6] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 2010. 5
- [7] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, 2020. 4
- [8] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *International Conference on Learning Representations*, 2022. 2
- [9] Akshita Gupta, Sanath Narayan, KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Ow-detr: Open-world detection transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 5
- [10] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, 2019. 3, 5
- [11] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 2, 3
- [12] Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. Lorahub: Efficient cross-task generalization via dynamic lora composition. *arXiv preprint arXiv:2307.13269*, 2023. 3
- [13] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *International Conference on Learning Representations*, 2023. 7
- [14] 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. *Proceedings of the National Academy of Sciences*, 2017. 2
- [15] Jahanzaib Latif, Chuangbai Xiao, Azhar Imran, and Shanshan Tu. Medical imaging using machine learning and deep learning algorithms: a review. In *International conference on computing, mathematics and engineering technologies*, 2019. 1
- [16] Chunyuan Li, Haotian Liu, Liunian Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping Jin, Houdong Hu, Zicheng Liu, Yong Jae Lee, et al. Elevater: A benchmark and toolkit for evaluating language-augmented visual models. *Advances in Neural Information Processing Systems*, 2022. 2
- [17] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2, 5, 12, 13
- [18] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *International Joint Conference on Natural Language Processing*, 2021. 3
- [19] 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 *Proceedings of the European Conference on Computer Vision*, 2014. 5, 12
- [20] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. In *International Conference on Learning Representations*, 2022. 1
- [21] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with

- grounded pre-training for open-set object detection. In *Proceedings of the European Conference on Computer Vision*, 2024. 1, 2
- [22] Yaoyao Liu, Bernt Schiele, Andrea Vedaldi, and Christian Rupprecht. Continual detection transformer for incremental object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 5
- [23] Haoyu Lu, Yuqi Huo, Guoxing Yang, Zhiwu Lu, Wei Zhan, Masayoshi Tomizuka, and Mingyu Ding. Uniadapter: Unified parameter-efficient transfer learning for cross-modal modeling. In *International Conference on Learning Representations*, 2024. 3
- [24] Imad Eddine Marouf, Subhankar Roy, Enzo Tartaglione, and Stéphane Lathuilière. Weighted ensemble models are strong continual learners. In *Proceedings of the European Conference on Computer Vision*, 2024. 8
- [25] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, 1989. 5
- [26] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional detr for fast training convergence. In *IEEE International Conference on Computer Vision*, 2021. 1
- [27] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Dilan Gorur, Razvan Pascanu, and Hassan Ghasemzadeh. Linear mode connectivity in multitask and continual learning. In *International Conference on Learning Representations*, 2021. 4
- [28] Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *Advances in Neural Information Processing Systems*, 2024. 7
- [29] Mark Palatucci, Dean Pomerleau, Geoffrey E Hinton, and Tom M Mitchell. Zero-shot learning with semantic output codes. *Advances in Neural Information Processing Systems*, 2009. 2
- [30] Jonas Pfeiffer, Sebastian Ruder, Ivan Vulić, and Edoardo Ponti. Modular deep learning. *Transactions on Machine Learning Research*, 2023. 2
- [31] Angelo Porrello, Lorenzo Bonicelli, Pietro Buzzega, Monica Millunzi, Simone Calderara, and Rita Cucchiara. A second-order perspective on model compositionality and incremental learning. In *International Conference on Learning Representations*, 2025. 3
- [32] Christian Poss, Olim Ibragimov, Anoshan Indreswaran, Nils Gutsche, Thomas Irrenhauser, Marco Prueglmeier, and Daniel Goehring. Application of open source deep neural networks for object detection in industrial environments. In *IEEE International Conference on Machine Learning and Applications*, 2018. 1
- [33] Yujia Qin, Cheng Qian, Jing Yi, Weize Chen, Yankai Lin, Xu Han, Zhiyuan Liu, Maosong Sun, and Jie Zhou. Exploring mode connectivity for pre-trained language models. In *Conference on Empirical Methods in Natural Language Processing*, 2022. 4
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. 1, 2
- [35] Apoorva Raghunandan, Pakala Raghav, HV Ravish Aradhya, et al. Object detection algorithms for video surveillance applications. In *International Conference on Communication and Signal Processing*, 2018. 1
- [36] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [37] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. In *Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. 2
- [38] Riccardo Salami, Pietro Buzzega, Matteo Mosconi, Jacopo Bonato, Luigi Sabetta, and Simone Calderara. Closed-form merging of parameter-efficient modules for federated continual learning. In *International Conference on Learning Representations*, 2025. 3
- [39] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. 13
- [40] James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim, Assaf Arbelle, Rameswar Panda, Rogerio Feris, and Zsolt Kira. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 3
- [41] Guido M Van de Ven, Tinne Tuytelaars, and Andreas S Tolias. Three types of incremental learning. *Nature Machine Intelligence*, 2022. 2, 3
- [42] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 1, 3
- [43] Xin Wang, Thomas E Huang, Trevor Darrell, Joseph E Gonzalez, and Fisher Yu. Frustratingly simple few-shot object detection. In *International Conference on Machine Learning*, 2020. 5
- [44] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *Proceedings of the European Conference on Computer Vision*, 2022. 3
- [45] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2, 3
- [46] Xun Wu, Shaohan Huang, and Furu Wei. Mixture of lora experts. In *International Conference on Learning Representations*, 2024. 3
- [47] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference

when merging models. *Advances in Neural Information Processing Systems*, 2024.

- [48] Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. In *International Conference on Learning Representations*, 2024. [3](#)
- [49] Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. Open-vocabulary detr with conditional matching. In *Proceedings of the European Conference on Computer Vision*, 2022. [2](#)
- [50] Jinghan Zhang, Junteng Liu, Junxian He, et al. Composing parameter-efficient modules with arithmetic operation. *Advances in Neural Information Processing Systems*, 2023. [3](#)
- [51] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations*, 2021. [1](#)

Supplementary Material Overview

This supplementary material complements the main manuscript by providing additional analyses, insights, and results on DitHub. Specifically, we delve deeper into efficiency considerations of our adaptation approach (App. A), highlighting blockwise adaptation and LoRA specializations. We introduce the novel ODinW-O benchmark for incremental learning with overlapping classes and domain shifts (App. B), and present insights on unlearning class-specific knowledge (App. D.2). Additionally, we clarify critical implementation specifics (App. C), present extensive robustness checks (App. D.1), and showcase qualitative examples (App. E) to illustrate the practical benefits of our methodology.

A. Further Insights on Efficiency

A.1. Blockwise Adaptation Analysis

While Grounding DINO consists of four main components — an image backbone for visual encoding, a textual backbone for language understanding, and an encoder-decoder pair for detection — we focus our adaptation on the latter two. This choice is motivated by the need for a more detection-centric adaptation and efficiency considerations, as the textual backbone, being a BERT Base model [3], accounts for more than half of Grounding DINO’s parameters. In Table B, we present results for adapting only the encoder (Figure Aa), only the decoder (Figure Ab), and both encoder and decoder (Figure Ac). DitHub employs the first approach — adapting only the encoder — as it delivers the best performance on both ZCOCO and Avg metrics.

A.2. Employing Full LoRA for Class-Specificity

While our methodology leverages a shared and unique B matrix for all tasks, Tab. A presents results obtained using class-specific pairs. In this variant of DitHub, both the A and B low-rank matrices of LoRA undergo the same specialized optimization process: an initial warmup phase followed by dedicated specialization. In other words, the training pipeline for matrix B mirrors that of matrix A described in Section 4.1. Applying our full set of strategies to both matrices yields a + 1.03 improvement in average performance across the tasks of ODinW-13, albeit at the cost of doubling the number of trainable parameters (Fig. 3). Prioritizing efficiency, we opt for a lightweight implementation that trades a marginal performance gain for significantly reduced memory requirements. Notably, ZCOCO results favor our efficient approach, as a single, unspecialized B matrix remains more aligned with the pretraining distribution, thereby mitigating degradation in zero-shot capabilities.

Method	ZCOCO	Avg
DitHub	47.01	62.19
DitHub (LoRA full)	46.66	63.22

Table A. LoRA full results on ODinW-13, with DitHub results taken from Table 1 for reference.

B. Additional Benchmark Details

This section provides additional details regarding the datasets utilized in our evaluation framework. We first outline the key characteristics of ODinW-13 and MS COCO, as these datasets constitute the foundation for evaluating the specialization and zero-shot capabilities of DitHub. Subsequently, we introduce and discuss the ODinW-O benchmark in greater depth, which we explicitly design to evaluate incremental learning performance under scenarios involving recurring categories across multiple tasks, accompanied by domain shifts. By leveraging these datasets, we establish a rigorous evaluation protocol to comprehensively assess model robustness in challenging real-world scenarios.

B.1. ODinW-13 and MS COCO

Following [2], we assess our model specialization using the Object Detection in the Wild (ODinW-13) [17] benchmark and the MS COCO [19] dataset to verify whether the zero-shot capabilities are preserved after the sequential fine-tuning on ODinW-13.

Unlike structured and standardized datasets such as MS COCO, ODinW-13 aggregates data from thirteen sources: *Aerial Maritime Drone* (Ae), *Aquarium* (Aq), *Cottontail Rabbit* (Co), *EgoHands* (Eg), *Mushrooms* (Mu), *Packages* (Pa), *Pascal VOC* (Pv), *Pistols* (Pi), *Pothole* (Po), *Raccoon* (Ra), *Shellfish* (Sh), *Thermal Dogs and People* (Th), *Vehicles* (Ve). This comprehensive collection covers rare object categories, diverse contextual conditions, and challenging domains. Consequently, the ODinW-13 benchmark offers an ideal scenario to evaluate the capability of DitHub to sequentially adapt to heterogeneous task distributions while effectively mitigating catastrophic forgetting on previously learned tasks.

Following the convention established by [2], we use the term ZCOCO to denote the zero-shot evaluation conducted on the MS COCO validation set comprising 5000 samples and consisting of annotations for 80 common object categories. The ZCOCO evaluation specifically measures the robustness and generalization capabilities of DitHub in zero-shot scenarios, after training on each successive task from ODinW-13.

B.2. ODinW-O

We introduce a new benchmark, ODinW-O (Overlapped), designed to further evaluate the class-incremental and spe-




Figure A. Blockwise analysis of adaptation strategies: (a) Adaptation of the encoder only (our chosen approach), (b) Adaptation of the decoder only and (c) Joint adaptation of both the decoder and encoder.

Method	ZCOCO	Avg	Ae	Aq	Co	Eg	Mu	Pa	Pv	Pi	Po	Ra	Sh	Th	Ve
DitHub	47.01	62.19	34.62	50.65	70.46	68.56	49.28	65.57	69.58	71.10	56.65	70.88	52.82	79.30	68.18
Only Dec	46.11	54.35	28.00	35.22	65.28	65.11	37.62	59.71	68.11	66.27	35.22	67.34	39.46	73.87	65.30
Dec/Enc	46.55	59.92	32.84	43.16	71.47	67.44	49.78	62.84	70.32	70.83	49.24	70.89	46.12	77.05	67.48

Table B. Performance comparison across datasets when adapting only the decoder ("Only Dec"), both encoder and decoder ("Dec/Enc"), and encoder only (DitHub).

cialization capabilities of DitHub. Derived from ODinW-35, ODinW-O retains only classes shared across at least one task (*overlapping classes*), while removing tasks with no shared classes. This filtering is necessary because classes that do not reappear in subsequent tasks present an ideal case in the DitHub framework. Specifically, their corresponding non-recurring class modules remain frozen, requiring no further adaptation. In contrast, ODinW-O eliminates this scenario by ensuring that all retained classes appear in multiple tasks. By enforcing strictly overlapping classes, ODinW-O provides a challenging benchmark for evaluating incremental learning, specifically tailored to assess the robustness of our method. Additionally, to rigorously evaluate the model robustness to domain shifts, we designed ODinW-O by explicitly selecting datasets that feature markedly different contexts (Fig. B), enabling clear insights into how effectively the model adapts to variations in visual appearance, environmental conditions, and imaging domains. We select the following datasets, already presented in App. B.1, and extracted only the overlapping classes:

- **Aerial Maritime Drone:** *boat* and *car*,
- **Pascal VOC:** *boat*, *car*, *dog*, and *person*
- **Thermal Dogs and People:** *dog* and *person*,
- **Vehicles:** *car* and *truck*,

Additionally, we introduce two further datasets from ODinW-35 to emphasize domain variation:

- **Hard Hat Workers:** *person*
- **Self Driving Cars:** *car* and *truck*

Despite ODinW-O enforcing class overlap and introducing domain shifts, our model consistently achieves state-of-the-art performance by effectively refining class knowledge while adapting across diverse contexts.



Figure B. Visualization of domain shifts for shared categories across different tasks within the ODinW-O benchmark.

C. Implementation Details

All experiments for DitHub were conducted on a single RTX A5000 GPU with 24GB of VRAM, requiring approximately 8 hours per training run with a batch size of 2. During training, our implementation in PyTorch efficiently utilizes around 10 GB of VRAM. We employed GroundingDINO-T (Swin-T) as an open-vocabulary object detector, pretrained on O365 [39], GoldG [17], and Cap4M [17] datasets. We performed training with random shuffling across three different seeds to ensure robustness.

For LoRA fine-tuning, we perform low-rank updates to the attention layers with a rank of 16, applying a learning rate of 1×10^{-3} and a weight decay of 1×10^{-2} . This approach enables efficient adaptation while significantly reducing computational overhead and memory storage for

Method	ZCOCO	Avg	Ae	Hw	Pv	Sd	Th	Ve
ZiRa	± 0.25	± 0.27	± 1.76	± 1.14	± 1.35	± 0.32	± 1.75	± 1.76
Ours	± 0.24	± 0.66	± 2.39	± 1.88	± 0.76	± 1.31	± 0.46	± 0.84

Table C. Standard deviations of the reproduced methods in Table 2.




Figure C. Matrix visualization of mAP degradation per class after removing class modules. Colors are normalized column-wise for readability; darker cells indicate greater performance drops.

each class-specialized module. Optimization is performed using AdamW. We rely solely on the contrastive classification loss and the localization loss of Grounding DINO, without introducing any additional loss functions. The full Grounding DINO model, after integrating 44 classes during training on ODinW-13, consists of approximately 173 million parameters, with 19.6 million LoRA parameters actively updated during training. After fine-tuning on ODinW-13, only 75.65 MB of LoRA parameters are retained for storage.

D. Additional Results

D.1. Standard Deviations

We report the standard deviations of the reproduced methods from Tab. 1 and Tab. 2 in Tab. D and C, respectively, computed over three runs with different random seeds.

D.2. Unlearning Additional

As discussed in Sec. 5.2, our modular approach enables the precise removal of specific knowledge from Grounding DINO by subtracting corresponding experts. Figure C illustrates this process: each column represents a target class's

mAP after unlearning different modules, with colors normalized column-wise to highlight relative drops. The darker diagonal pattern indicates that removing a class module leads to the most significant performance drop for its corresponding class. This confirms that our modules effectively encapsulate specialized knowledge, as unlearning them selectively impacts their respective classes.

E. Qualitatives

In Fig. D, we compare zero-shot Grounding DINO (left) with ZiRA (center) and DitHub (right). Our method excels at detecting challenging objects, including small ones that other models miss. Additionally, DitHub avoids a false positive (ZiRA) and a false negative (Grounding DINO), demonstrating more reliable predictions. It also correctly identifies both cars (last row) despite occlusions and distractors. This improved stability stems from per-class specialization, which mitigates hallucinations and ensures smoother adaptation to domain shifts.

Shots	Method	ZCOCO	Avg	Ae	Aq	Co	Eg	Mu	Pa	Pv	Pi	Po	Ra	Sh	Th	Ve
1	ZiRa	± 0.90	± 2.05	± 1.49	± 0.17	± 2.16	± 0.98	± 0.79	± 1.63	± 1.33	± 7.91	± 3.29	± 12.17	± 4.45	± 3.07	± 3.25
	DitHub	± 0.60	± 0.81	± 0.95	± 0.57	± 2.66	± 1.58	± 3.72	± 0.00	± 1.53	± 1.10	± 2.90	± 7.03	± 3.68	± 4.66	± 2.46
5	ZiRa	± 0.23	± 0.43	± 3.76	± 2.15	± 1.40	± 1.08	± 3.54	± 3.74	± 1.54	± 4.26	± 3.73	± 4.13	± 2.03	± 2.02	± 4.16
	DitHub	± 0.42	± 1.27	± 2.52	± 0.70	± 2.42	± 6.74	± 0.76	± 3.27	± 0.17	± 5.49	± 4.79	± 2.22	± 4.29	± 1.81	± 1.54
10	ZiRa	± 0.28	± 0.66	± 0.44	± 2.64	± 2.77	± 0.52	± 6.43	± 0.97	± 1.04	± 2.17	± 2.04	± 3.90	± 1.21	± 3.71	± 0.87
	DitHub	± 0.20	± 0.54	± 2.26	± 0.72	± 4.23	± 3.53	± 1.35	± 3.19	± 0.30	± 2.74	± 0.34	± 2.99	± 1.01	± 1.15	± 1.27
Full	ZiRA	± 0.44	± 0.68	± 2.73	± 1.17	± 1.39	± 4.77	± 7.26	± 2.19	± 1.33	± 0.33	± 4.14	± 3.57	± 2.44	± 2.82	± 3.13
	DitHub	± 0.16	± 0.44	± 2.39	± 1.00	± 2.83	± 0.48	± 1.34	± 0.60	± 0.38	± 1.68	± 0.76	± 1.57	± 2.87	± 0.69	± 1.01

Table D. Standard deviations of the reproduced methods in Table 1.

Grounding DINO

ZiRa

DitHub




Figure D. Qualitative comparison of Grounding DINO (left), ZiRA (center), and DitHub (right).