

Category Adaptation Meets Projected Distillation in Generalized Continual Category Discovery

Grzegorz Rypeś^{*1,2}, Daniel Marczak^{1,2}, Sebastian Cygert^{1,3},
Tomasz Trzciński^{1,2,4}, and Bartłomiej Twardowski^{1,5,6}

¹ IDEAS NCBR

² Warsaw University of Technology

³ Gdańsk University of Technology

⁴ Tooploox

⁵ Autonomous University of Barcelona

⁶ Computer Vision Center

Abstract. Generalized Continual Category Discovery (GCCD) tackles learning from sequentially arriving, partially labeled datasets while uncovering new categories. Traditional methods depend on feature distillation to prevent forgetting the old knowledge. However, this strategy restricts the model’s ability to adapt and effectively distinguish new categories. To address this, we introduce a novel technique integrating a learnable projector with feature distillation, thus enhancing model adaptability without sacrificing past knowledge. The resulting distribution shift of the previously learned categories is mitigated with the auxiliary category adaptation network. We demonstrate that while each component offers modest benefits individually, their combination – dubbed CAMP (Category Adaptation Meets Projected distillation) – significantly improves the balance between learning new information and retaining old. CAMP exhibits superior performance across several GCCD and Class Incremental Learning scenarios. The code is available on Github.

1 Introduction

Traditional machine learning models provide remarkable performances across various applications. However, they typically operate within the closed-world scenario, which assumes that all the tasks are well-defined and knowledge necessary to solve them is given upfront. Unfortunately, those assumptions are far from real-life applications, where models encounter dynamic streams of ever-changing data that evolves and may contain semantics the models have never encountered before. This variability severely impacts their performance and robustness. To study these problems from a unified perspective, many recent works combine Generalized Category Discovery [36] (GCD) setting with Continual Learning [25, 27] (CL). Resulting combinations [20, 39, 44, 45] (to which we refer as Generalized Continual Category Discovery - GCCD), consider a sequence of tasks, in which the method

^{*} Corresponding author, email: grzegorz.rypes.dokt@pw.edu.pl

has to continuously handle new data, discover novel categories and maintain good performance on previous tasks by overcoming catastrophic forgetting [12].

A popular trend in the field of discovering categories is to regularize the model with feature distillation [18, 20, 31] (FD) to alleviate forgetting. However, as pointed out by [13], it decreases the plasticity of the model by restricting changes (drift) of the categories’ distributions in the feature space. In this work, we show that in order to enable learning robust representations of new data, it is necessary to allow distributions of past categories to drift. Drawing inspiration from recent studies in representation learning for enhanced transferability [3, 34], we incorporate projected representation in a distillation-based regularization for continual learning. Compared to standard FD, this method results in improved plasticity. However, it increases forgetting as it lets old categories drift even more.

To address this issue, we represent categories as centroids in the latent space of the feature extractor and train an auxiliary network to predict the drift of centroids, which we call category adaptation. This enables the model to recover the actual positions of old clusters and improve the overall accuracy of the Nearest Centroid Classifier (NCC). Similarly, works [17, 42] analyze the semantic drift of old categories in closed-world, Continual Learning scenarios. Here, our key observation is that, while centroid adaptation demonstrates overall solid performance, its true effectiveness is realized only through combining it with feature distillation using a projection network, as depicted in Fig. 1. That allows the technique to excel. Building upon these findings, we propose to leverage this combined effect for GCCD by introducing a method called CAMP – Category Adaptation Meets Projected distillation. CAMP presents state-of-the-art performance across GCCD scenarios and exemplar-free Class Incremental Learning (CIL).

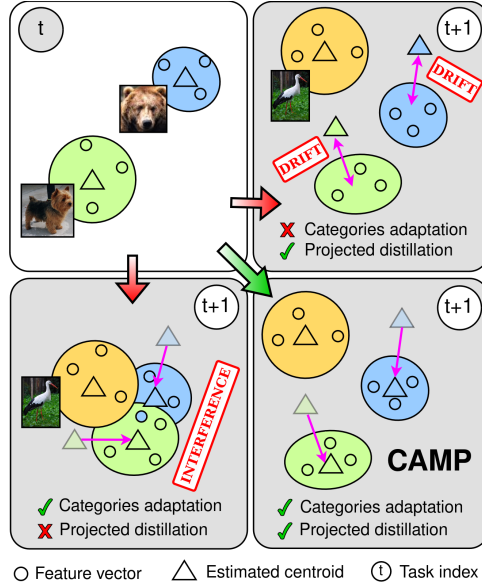


Fig. 1: While knowledge distillation through a projector and category adaption are decent on their own, CAMP combines them to achieve great results without the need of exemplars.

To summarize, the main contributions of our work are as follows. **1:** We study the interplay between knowledge distillation and category adaptation for the problem of Generalized Continual Category Discovery (GCCD). We demonstrate that they address distinct or even opposing challenges when considered separately. However, when combined, they complement each other significantly, reduce

forgetting, and improve overall results. **2:** Based on those insights, we propose a novel method for generalized continual category discovery dubbed CAMP. Built on GCD [36], CAMP improves it for continual learning by adapting categories between tasks and performing knowledge distillation with two separate projectors. **3:** We evaluate our method on numerous GCCD and Class Incremental Learning (CIL) datasets. With a series of experiments, we show that CAMP achieves state-of-the-art results with and without exemplars in different scenarios.

2 Related work

Continual learning (CL) is a setting where an agent learns a sequence of tasks with access to the data from the current task only. The goal is to achieve high performance on all the tasks from the sequence. Regularization-based approaches, such as EWC [22] or MAS [1] penalize changes to important neural network parameters. Alternatively, it is possible to regularize neuron activations as in LwF [25], DER [41] or PODNet [9] by using distillation techniques. However, even with feature distillation, the features from old classes will change, which causes catastrophic forgetting. [42] tries to predict these changes by approximating the drift of old classes based on the drift of current task data. [17] learns a feature adaptation network that translates old features into those learned on new tasks. In our work, we show that such an approach may not necessarily work well when going beyond the closed-world assumption under which CL usually operates.

Novel Class Discovery (NCL) [11, 14–16, 47, 48] aims to discover unknown classes from unlabeled data, thus it relaxes the closed-world assumption. NCL typically assumes that the class space of unlabeled data is completely disjoint with the labeled (training) set. This assumption was removed by Generalized Category Discovery (GCD) [36, 38], where a training dataset consists of two parts: known classes where only a fraction of samples are labeled and novel classes which are unlabeled. The goal is to categorize all images from known and novel classes, without any additional information regarding their novelty. Typical solutions in this area include unsupervised clustering [16, 46], contrastive learning [36, 38, 46], pseudo-labeling [38] or pairwise similarities [43].

Continual classes discovery combines the NCL and CL fields. In [32], learning of the novel classes is framed incrementally: that is, in the first stage, the model learns from supervised data, and in the following stages, novel classes are learned without any labels. A similar setting was explored in [18] with the solution that uses pseudo-latent supervision, feature distillation, and a mutual information-based regularizer. However, their setting assumes no class overlap between the initial and incremental training stages, which was later relaxed by various to which we refer as Generalized Continual Category Discovery - GCCD [20, 39, 44, 45]. Grow and Merge (GM) [44] keeps two models: a static one that maintains knowledge of old classes and a dynamic one trained using self-supervision to discover novel classes. Both models are then merged into one when the new task arrives. In IGCD [45] a non-parametric classifier and a small buffer of exemplars from previous tasks to decrease forgetting. PA [20] combines

exemplars and proxy anchors [21] known from deep metric learning. Contrary to previous approaches [20, 44, 45] we focus on the most challenging exemplar-free setting. MetaGCD [39] needs to store the data from the first task, which we avoid, by proposing a purely continual model.

3 Method

3.1 Problem Setting

In Generalized Continual Category Discovery models are trained sequentially on partially labeled tasks containing known and novel categories. Formally, we consider a scenario where a dataset D is partitioned into disjoint subsets, each represented by a task: $D = T^1 \cup T^2 \cup \dots \cup T^N$, with N denoting the total number of tasks. Each task T^t comprises two subsets: $T^t = T_L^t \cup T_U^t$, wherein $T_L^t = \{(x, y) \mid x \in \mathcal{X}_L^t, y \in \mathcal{Y}^t\}$ denotes the labeled set of known classes and $T_U^t = \{x \mid x \in \mathcal{X}_U^t\}$ encompasses unlabeled data samples from both known and novel classes.

Within the GCCD framework, each method engages in sequential learning from a stream of tasks. We assume that, at task t , each method comprises a feature extractor $\mathcal{F}^t : \mathcal{X}^t \rightarrow \mathcal{Z}^t$ and a classifier $C^t : \mathcal{Z}^t \rightarrow \mathcal{Y}^t$. Here, \mathcal{Z}^t signifies the latent space generated by the feature extractor trained on the task t . The GCCD model aims to learn to recognize categories of a given task T^t while also maintaining the ability to recognize the categories from previous tasks $T^1 \cup T^2 \cup \dots \cup T^{t-1}$. The main challenge is to prevent catastrophic forgetting on previous tasks while maintaining high plasticity that allows a method to learn new categories on new data.

3.2 Motivation

In this section, we demonstrate the motivation for our method on a simple experiment performed on ten categories of the CUB200 dataset split into two tasks. We train a neural network with its latent space bottlenecked to two dimensions (see Supplementary Material B for experimental details). In Fig. 2, we present a visualization of the latent space after training on the second task. We observe that naive fine-tuning of GCD (left) causes catastrophic forgetting and results in poor performance on the first task. Training with feature distillation (middle) [18, 20, 31] to some extent alleviates the forgetting and makes the drift of the centroids partially reversible via centroid adaptation (black arrows). However, feature distillation introduces rigid regularization that limits the model’s ability to acquire new knowledge and results in low performance on the second task. These results highlight the need for a mechanism that relaxes anti-forgetting regularization constraints while alleviating forgetting. Our proposed CAMP method (right) is able to address these issues by introducing a learnable projector, which makes the regularization less rigid and allows the distribution to drift more, resulting in better accuracy on the second task. At the same time, it causes the

drift to be easily predictable, thus enabling our centroid adaptation to estimate ground truth positions of centroids, leading to great accuracy on the first task. Existing feature adaptation methods [17, 42] do not utilize projected distillation. Therefore, their drift is more challenging to predict (middle). Additionally, our method is more memory efficient than [17] as we store only a single centroid instead of many features per class and we don't require exemplars.

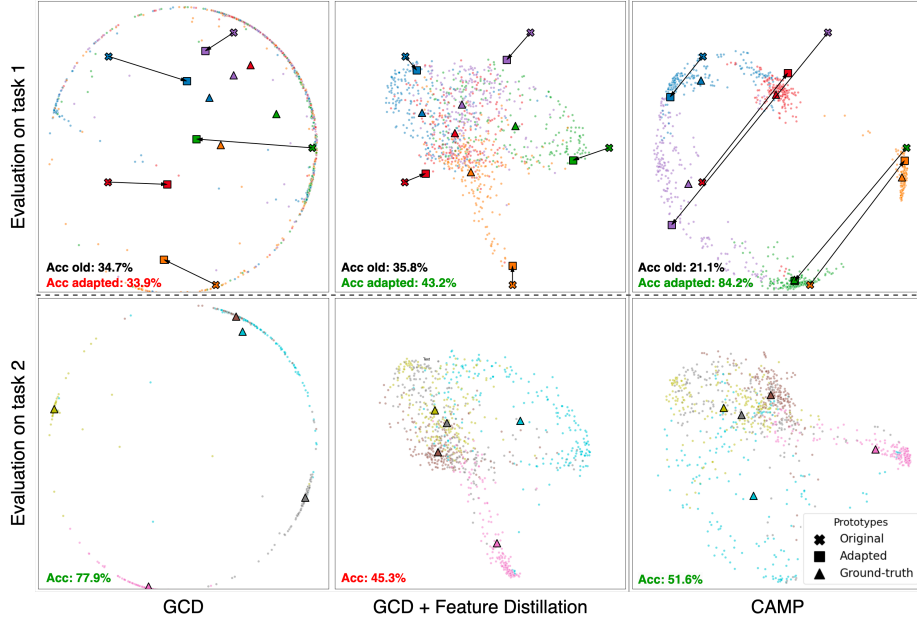


Fig. 2: CAMP utilizes a projected knowledge distillation, resulting in a predictable latent space drift. The drift is revertible via centroid adaptation (black arrows), maintaining high performance on the first task. GCD and GCD with feature distillation fail to prevent forgetting, resulting in a drift that is difficult to predict. This decreases their performance. We report the nearest centroid classification accuracy using stored (Acc old) and adapted centroids (Acc adapted) after training on the second task.

3.3 CAMP

The training procedure for CAMP within each task is divided into three consecutive phases: (1) feature extractor training, (2) clustering on a new task, and (3) adaptation of centroids representing categories of past classes. An overview is depicted in Fig. 3. The first phase focuses on training \mathcal{F}^t to transform images into meaningful features that effectively distinguish between different classes. The second phase is dedicated to identifying novel classes. The third phase, centroid adaptation, is responsible for updating the centroids of previous tasks by esti-

minating the drift caused by updating the feature extractor. Detailed explanations of each phase will follow in the subsequent three subsections.

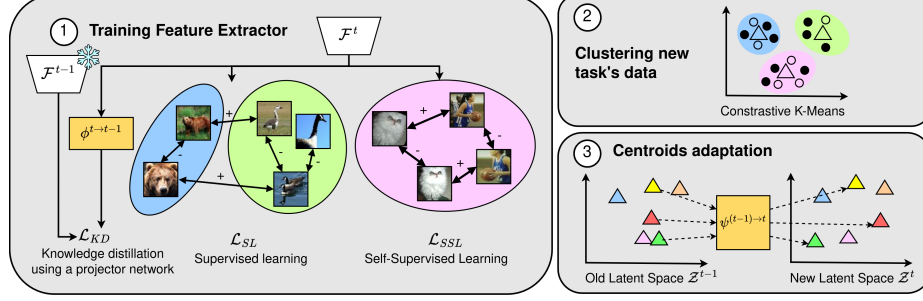


Fig. 3: The training procedure of CAMP consists of three stages: (1) We train the feature extractor in a semi-supervised manner using distillation through a learnable projector. (2) We obtain centroids of known and novel classes in the current task using contrastive K-Means algorithm. (3) We update memorized centroids of old categories to alleviate forgetting.

3.4 Phase 1: Feature extractor training

Our method for training feature extractors consists of three components: self-supervised learning, supervised learning, and knowledge distillation. The first two components allow the use of labeled and unlabeled data for the training feature extractor, while knowledge distillation is used for regularization and preventing forgetting. All are depicted in Fig. 3, left rounded rectangle.

Representation learning To learn the representations of the data from the new task, we follow SimGCD [38] and combine self-supervised and supervised learning. We use self-supervision on all data from known and novel categories T_t^U . We combine SimCLR [6] loss and cross-entropy loss between predictions and pseudo-labels: $\mathcal{L}_{SSL} = \mathcal{L}_{SimCLR} + \mathcal{L}_{pseudo}$. For labeled data T_t^L we use supervised learning loss that combines SupCon [19] loss and cross-entropy loss between predicted and ground-truth labels: $\mathcal{L}_{SL} = \mathcal{L}_{SupCon} + \mathcal{L}_{CE}$. A detailed formal description of each loss component is provided in Supplementary Material A.

Knowledge distillation We use distillation to alleviate forgetting in continual learning scenarios. Following state-of-the-art exemplar-free continual representation learning methods [10, 13], we use a learnable MLP projector $\phi^{t \rightarrow t-1} : \mathcal{Z}^t \rightarrow \mathcal{Z}^{t-1}$ that maps the features learned on the current task back to the old latent space. In our case, this approach improves the trade-off between plasticity and stability. Additionally, it ensures there is a mapping between old $(t-1)$ and new (t) latent space, which can be later inverted during category adaptation.

Following distillation-based continual learning methods [10, 13, 25], we freeze previously trained feature extractor \mathcal{F}^{t-1} and regularize currently trained feature extractor \mathcal{F}^t with the outputs of \mathcal{F}^{t-1} :

$$\mathcal{L}_{KD} = \sum_{i \in B} \|\phi^{t \rightarrow t-1}(\mathcal{F}^t(x_i)) - \mathcal{F}^{t-1}(x_i)\|^2. \quad (1)$$

where B is a batch of data. This form of distillation does not require labels. Therefore, all the data from T_U^t can be used for regularization.

The final loss function for feature extractor training is defined as follows:

$$\mathcal{L} = (1 - \alpha)((1 - \beta)\mathcal{L}_{SSL} + \beta\mathcal{L}_{SL}) + \alpha\mathcal{L}_{KD}, \quad (2)$$

where $\alpha, \beta \in [0, 1]$ are hyperparameters defining contribution of regularization and supervision respectively.

3.5 Phase 2: Clustering task’s data

The purpose of clustering the new task’s data T^t is to find a centroid of each known and novel class in the learned representation space of the feature extractor \mathcal{F}^t . Such centroid is the class representation in latent space, later used for classification with the Nearest Centroid Classifier (NCC) during inference time. In order to find class centroids for CAMP, we use a semi-supervised k-means algorithm introduced in [36]. We use labeled data samples of known classes (T_L^t) to initialize known classes’ centroids. We obtain the set of remaining centroids (equal to the number of novel classes) from the unlabeled data T_U^t using the k-means++ [2] algorithm, with the constraint that they must be close to the centroids of the labeled ones from T_L^t . To determine the number of classes in a task, we utilize a popular elbow method as in [16, 36]. Thus, for different numbers of K-Means clusters, we measure clustering accuracy on the labeled data and set the number of clusters to the one that yields the highest accuracy.

3.6 Phase 3: Centroid adaptation

Although regularization methods decrease forgetting, they still allow for a drift of ground truth representations of old classes (as discussed in Sec. 3.2). To recover from this issue, we propose to train auxiliary centroid adaptation network $\psi^{(t-1) \rightarrow t}$ to predict the drift of data representations from the previous task’s latent space ($t-1$) to the current one (t). To train it, we use all the data available in the current task. Exemplars can also be used for this purpose but are not necessary. We examine the architecture of $\psi^{(t-1) \rightarrow t}$ in Sec. 4.3. We train $\psi^{(t-1) \rightarrow t}$ by minimizing the MSE loss:

$$\mathcal{L}_{PA} = \sum_{i \in B} \|\mathcal{F}^t(x_i) - \psi^{(t-1) \rightarrow t}(\mathcal{F}^{t-1}(x_i))\|^2. \quad (3)$$

When $\psi^{(t-1) \rightarrow t}$ is trained, we use it to calculate the updated centroids $p_i^t = \psi^{(t-1) \rightarrow t}(p_i^{t-1})$ where p_i^t is a centroid of class i calculated after task t . This is presented in Fig. 1, lower right.

4 Experiments

4.1 Experimental setup

We evaluate all methods on five datasets: CIFAR100 [24], Stanford Cars [23], CUB200 [37], FGVC Aircraft [26] and DomainNet [28]. We split datasets into five equal tasks, with the exception of Stanford Cars (4 tasks) and DomainNet (6 tasks). Following [33] we set a different domain for each task of DomainNet, to simulate domain shifts. Following [36, 38], for each dataset, we set the ratio of known to novel classes as 4:1. For known classes, we set the ratio of labeled data samples to unlabeled data as 1:1. As the feature extractor \mathcal{F} for all methods we use ViT-small [8] model pretrained with DINO [4] on ImageNet [7]. Following GCD [36] we freeze the first 11 blocks for all methods. We train all methods for 100 epochs in each task using AdamW optimizer with a starting learning rate of 0.1^3 and cosine annealing scheduling. The exception is PA method, for which we set it to 0.1^4 but increase the learning rate for proxies 100 times following the original work. As the distiller ϕ we utilize a 2-layer MLP network consisting of 384 neurons and ReLU activation. As the adapter ψ , we utilize a linear layer consisting of 384 neurons. For CAMP without exemplars, we set α to 0.5; for CAMP with exemplars, we use $\alpha = 0.1$.

Selection of baselines For the simplest baseline, we fine-tune GCD [36] method and couple it with popular CL regularization-based methods: (1) weight regularization method - EWC [22] (GCD+EWC); (2) feature distillation method [13, 25] (GCD+FD). In GCD+FD we distill features $\mathcal{F}(x)$. Moreover, we combine GCD with experience replay [30] (GCD+ER), where for incremental tasks, we add a buffer that stores random images of past known classes. For all the above methods, we perform semi-supervised k-means [36] algorithm to find centroids and utilize the NCC classifier for classification. Additionally, we use SimGCD [38], which improves GCD by utilizing a learnable classifier and adding additional cross-entropy-based losses for training it. We also test MetaGCD [39] that utilizes meta-learning and PA [20] method, which trains the feature extractor using supervised proxy anchor loss, maintains a buffer of exemplars, and uses feature distillation to alleviate forgetting. Lastly, we evaluate IGCD [45] that utilizes a density-based support set selection and replay buffer to mitigate forgetting. We describe all methods in detail in the Supplementary Material C.

Evaluation protocol We assess the performance of each method upon completion of the final task using the test set for each task. Our primary evaluation metrics encompass final task-agnostic (TAg) accuracy [5] (where the model does not know task-id during the inference), average forgetting [5] after last task and plasticity defined as $\frac{1}{N-1} \sum_{n=2}^N A_n$, where A_n denotes accuracy measured on n-th task after training on n-th task. Following [36, 38], we calculate them separately for known, novel, and all classes to get more insight into the performance of methods. We will release the code upon the acceptance of the manuscript.

Table 1: CAMP achieves state-of-the-art results in most of the scenarios. We report known, novel, and all accuracy (%) calculated over all classes after the last incremental learning step. We assume that task id is not known during the inference.

Method	Exemplars per class	CIFAR100			Stanford Cars			CUB200			FGVC Aircraft			DomainNet		
		All	Known	Novel	All	Known	Novel	All	Known	Novel	All	Known	Novel	All	Known	Novel
GCD [36]	0	26.7	26.6	27.2	20.6	23.8	5.8	20.6	22.5	13.3	16.8	20.0	4.4	24.7	23.7	28.5
GCD+EWC [22]	0	26.9	28.7	19.7	30.4	35.3	8.7	50.3	53.3	38.1	25.3	28.2	13.8	33.1	37.5	16.8
GCD+FD	0	36.6	40.3	21.6	27.9	31.5	11.8	42.0	45.7	27.2	28.2	30.6	18.6	32.7	34.4	39.7
MetaGCD [39]	0	35.8	40.0	19.1	23.6	25.6	14.4	42.0	44.6	31.4	28.6	31.0	19.0	30.8	34.2	37.0
SimGCD [38]	0	31.2	31.9	28.5	21.3	30.1	23.9	35.7	36.6	31.8	22.0	22.3	20.8	36.5	36.4	35.0
CAMP	0	52.1	55.7	37.8	48.8	52.8	31.0	58.9	62.6	44.2	39.9	42.0	31.5	36.7	39.2	29.7
GCD+ER [30]	5	31.7	32.9	27.0	32.4	38.1	6.5	36.6	42.4	13.7	36.6	42.4	13.7	34.2	38.4	34.2
PA [20]	5	35.9	40.8	16.2	42.8	48.1	18.7	50.1	55.7	27.4	40.5	43.8	27.6	36.0	43.1	12.4
MetaGCD [39]	5	38.9	43.6	20.0	34.5	38.9	14.5	48.2	52.0	32.9	37.8	41.3	24.0	34.8	38.5	22.1
IGCD [45]	5	40.6	43.6	28.8	37.8	41.3	22.4	50.8	54.8	34.8	34.5	36.4	26.9	45.9	35.8	43.4
CAMP	5	52.7	56.1	39.1	50.3	55.0	29.1	61.1	65.1	45.0	40.3	44.6	23.0	45.1	48.8	29.6
GCD+ER [30]	20	42.5	46.3	27.3	46.3	54.0	11.5	52.0	59.9	20.1	48.9	56.5	17.5	46.4	53.7	20.4
PA [20]	20	37.6	43.1	15.7	46.7	53.0	18.7	53.1	58.7	30.4	45.6	50.5	26.0	38.2	42.0	25.3
MetaGCD [39]	20	43.6	47.3	28.6	47.9	51.5	31.8	55.8	56.9	42.2	46.5	49.7	33.9	42.4	47.6	25.8
IGCD [45]	20	51.8	57.3	29.6	57.6	63.9	29.7	64.0	69.0	43.4	49.5	52.6	37.2	51.1	58.3	22.0
CAMP	20	58.0	63.6	35.7	60.5	66.5	33.6	66.1	71.6	43.9	50.8	56.7	27.0	56.9	64.3	26.6

4.2 Comparison to baselines

We compare CAMP to other methods in terms of accuracy in Tab. 1. CAMP outperforms existing methods by a large margin in most scenarios, achieving the best results on CIFAR100, Stanford Cars, and CUB200 with and without exemplars, *e.g.* on CUB200, CAMP with no exemplars achieves 8.6%, 9.3%, 6.1% points better all, known and novel accuracy than the second-best method GCD+EWC. However, results are not that consistent on FGVC Aircraft and DomainNet. On DomainNet GCD+FD method achieves 10.0% points better novel accuracy than CAMP, yet still, our method has 4.0% points better all accuracy. We can also study the impact of exemplars. Methods that utilize exemplars outperform exemplar-free competitors at the cost of additional memory complexity. Exemplars drastically improve GCD method, on CUB200 they boost its all accuracy by 16.0% and 31.4% points with 5 and 20 exemplars, respectively.

To better examine how the accuracy changes in incremental steps, we plot averages of all accuracy after each task in Fig. 4 (left). CAMP achieves the highest accuracy from the second task, and the results are consistent with and without exemplars. Next, we analyze the plasticity and forgetting of all methods in Fig. 4 (right). GCD is the most plastic method. However, it has the highest forgetting as it has no form of regularization. On the other hand, IGCD is the most stable as it achieves the lowest forgetting. CAMP has good plasticity and forgetting, allowing it to achieve state-of-the-art results. We can also observe that exemplars prevent methods like GCD, IGCD, PA, and CAMP from forgetting. That is intuitive as the replay buffer allows them to recall the past categories’ data. Interestingly, each method forms its own cluster in presented plots, informing that they have different but specific characteristics for continual learning scenarios.

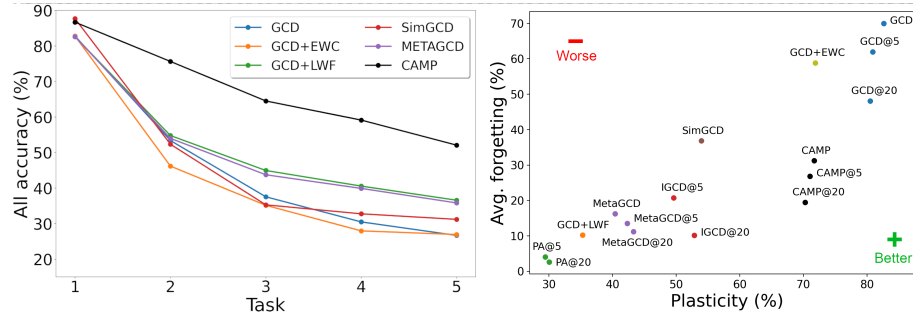


Fig. 4: Results of GCCD methods on CIFAR100 during the continual learning session of five tasks. We report average accuracy after each task (**left**); plasticity and forgetting (**right**). CAMP achieves better accuracy and than competitors.

Class Incremental Learning results Class incremental learning [27] is a special case of GCCD, where each data sample is labeled, and all classes are known. We test CAMP against exemplar-free state-of-the-art methods and train the feature extractor (ResNet-18) from scratch. We train methods for N equal tasks on CIFAR100, Domainnet, and Imagenet-Subset. For training CAMP we remove the self-supervised component from the loss function as all data is labeled. For baselines, we utilize popular exemplar-free CIL methods using their implementations in well-established benchmarks (FACIL [27], PyCIL [49]) and set their hyperparameters to default. We implement CAMP in FACIL and will publish the code upon the acceptance of the manuscript. For details, please refer to Supplementary Material B. We report average incremental accuracy [27], which is the average of task agnostic accuracies measured after each task on tasks seen so far. We provide results in Tab. 2.

CAMP outperforms baselines in terms of average incremental accuracy. On CIFAR100 it is better than the second best method - FeTrIL [29] by 6.5%, 10.4%, 4.3% for 5, 10, 20 tasks respectively. FeTrIL freezes the model after the first task, what decreases its ability to adapt to new data. CAMP does not, thus allowing for better plasticity. Results are consistent on DomainNet and ImageNet-Subset. These results show that CAMP translates well to Class Incremental Learning scenario, therefore combining projected distillation and category adaptation is a promising technique in the field of Continual Learning.

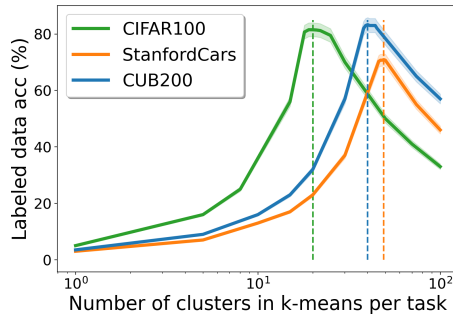
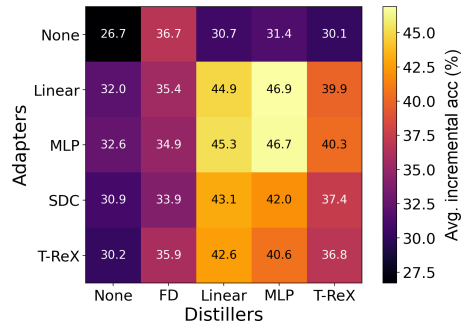
4.3 Method Analysis

How do distillers and adapters affect performance? We verify the influence of different network architectures of adapters and distillers on the accuracy achieved by CAMP. For this purpose, we evaluate three types of adapter networks ψ : linear, two connected layers with batch-norm and ReLU activation after the first layer (MLP) as in [6], and three layers following [35] (T-ReX). We also test a semantic drift compensation [42] (SDC) method to compare adaptation with different, non-learnable methods. It utilizes a vector field to estimate the drift

Table 2: Exemplar-free class incremental results for equal task split. We train all the methods from scratch and report task agnostic average incremental accuracy (%).

CIL Method	CIFAR-100			DomainNet			ImageNet-Subset		
	$N=5$	$N=10$	$N=20$	$N=6$	$N=12$	$N=24$	$N=5$	$N=10$	$N=20$
Fine-tune	40.6	26.4	17.1	27.7	17.9	14.8	41.6	27.4	18.6
GCD [36]	40.1	25.2	16.5	28.2	18.2	15.3	42.0	26.1	18.3
EWC [22]	52.9	37.8	21.0	30.0	19.2	15.7	43.7	29.8	19.1
LwF [25]	49.4	47.0	38.5	30.5	20.9	15.1	54.4	32.3	33.7
PASS [51]	62.2	52.0	41.4	33.6	25.9	14.9	63.6	53.9	36.0
IL2A [50]	-	43.5	28.3	31.3	20.7	18.2	-	-	-
SSRE [52]	57.0	44.2	32.1	34.9	33.2	24.0	56.9	45.0	32.9
FeTrIL [29]	58.5	46.3	38.7	36.8	33.5	27.5	62.9	58.7	43.2
CAMP	65.0	56.7	43.0	43.6	36.6	27.9	73.1	59.9	50.2
Joint	71.4			65.1	63.7	69.3	81.5		

of prototypes in the latent space, and here, we use it to adapt centroids. As a distiller network ϕ , we utilize popular feature distillation [18, 20, 31] technique and three networks of the same architecture as for adapters (Linear, MLP, T-ReX).

**Fig. 5:** Results for estimation of number of categories in tasks. We utilize an elbow method following vanilla GCD method. Accuracy calculated on labeled data peaks around ground truth number of clusters.**Fig. 6:** Performance of CAMP for different adapting (ψ) and distilling (ϕ) networks. It achieves the best results when MLP distiller is combined with Linear or MLP adapter.

We provide results for CIFAR100 in Fig. 6. It is best to combine MLP distiller and linear adapter, which justifies the design choice for CAMP. Utilizing only MLP distiller improves base results by 4.7% points. However, combining MLP and linear adapter improves results by 20.2%, which is a significant gain. We also achieve 12% better accuracy than combination inspired by [17] (MLP adapter and FD distiller). This shows that combining projected distillation with adaptation is crucial for achieving good performance. Interestingly, all adaptation methods

improve no-adaptation results if FD distiller is not used. That suggests that FD changes old representation in a harder-to-predict way than other distillers.

Adaptation with feature distillation or projected distillation? We examine the impact of Feature Distillation (FD) and distillation with an MLP projector on centroid adaptation for α parameter value in range $[0.03, 0.99]$ (plasticity-stability trade-off). We present the results in Fig. 7. We can see that without adaptation, FD imposes more substantial constraints on the drift of old centroids than MLP, as the distance between real and memorized centroids is lower (right). However, after adapting centroids this distance is much lower for MLP, proving that MLP facilitates drift prediction. Moreover, the adaptation method does not work well for FD with $\alpha \geq 0.4$, as the rigid features regularization minimizes the gain from centroid adaptation. The accuracy when using MLP is much better than FD, which proves the design choice of the distilling method.

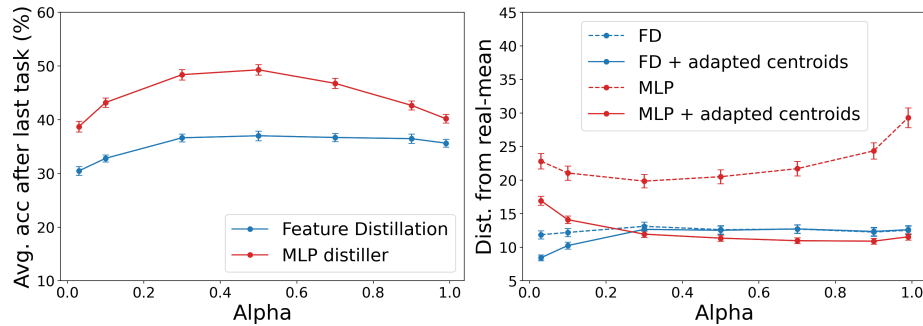


Fig. 7: MLP projector relaxes drifts of class distributions (measured with distance from memorized to ground-truth class centroids) compared to feature distillation. Additionally, it enables better adaptation of centroids, leading to increased accuracy.

Ablation study We perform ablation study on CUB200 and FGVC Aircraft datasets. We present results in Tab. 3. We test CAMP without: \mathcal{L}_{SL} , \mathcal{L}_{SSL} , knowledge distillation through the projector loss (\mathcal{L}_{KD}) and without centroid adaptation. We can see that having all of these elements is crucial for obtaining good accuracy on known and novel classes (62.6% and 44.2% respectively on CUB200). Interestingly, without the \mathcal{L}_{SSL} CAMP achieves 2.0% points better accuracy for known classes but 2.8% points lower accuracy for novel classes.

Estimating number of categories in a task. We estimate the number of novel categories in each task using the elbow method. We present results in Fig. 5. It is clear that accuracy on the labeled dataset peaks when the number of clusters for k-means algorithm is equal to the ground truth number of categories.

Does category adaptation improve position of centroids? After each task, we measure the distance between ground truth and estimated centroids with and without our category adaptation method. We present results in Fig. 8. Interestingly, we can see that novel classes tend to be further away from their

Table 3: Ablation study of CAMP.

\mathcal{L}_{SL}	\mathcal{L}_{SSL}	\mathcal{L}_{KD}	Cent. adapt.	CUB200		FGVCAircraft	
				Known	Novel	Known	Novel
\times	\checkmark	\checkmark	\checkmark	29.9 \pm 1.6	12.1 \pm 1.1	19.0 \pm 0.9	18.8 \pm 1.9
\checkmark	\times	\checkmark	\checkmark	64.6 \pm 2.3	41.4 \pm 3.5	43.5 \pm 1.8	16.6 \pm 2.0
\checkmark	\checkmark	\times	\checkmark	30.9 \pm 2.6	32.9 \pm 2.6	24.7 \pm 1.2	19.8 \pm 2.2
\checkmark	\checkmark	\checkmark	\times	36.5 \pm 1.9	28.7 \pm 1.7	29.6 \pm 1.7	24.2 \pm 2.8
\checkmark	\checkmark	\checkmark	\checkmark	62.6 \pm 2.4	44.2 \pm 3.7	42.0 \pm 1.9	31.5 \pm 3.5

estimated positions than known ones. Our centroid adaptation method consistently improves estimations of real centroids, *e.g.* after the last incremental step it decreases mismatch between positions of known classes by 44.8% and novel classes by 28.8% on CUB200.

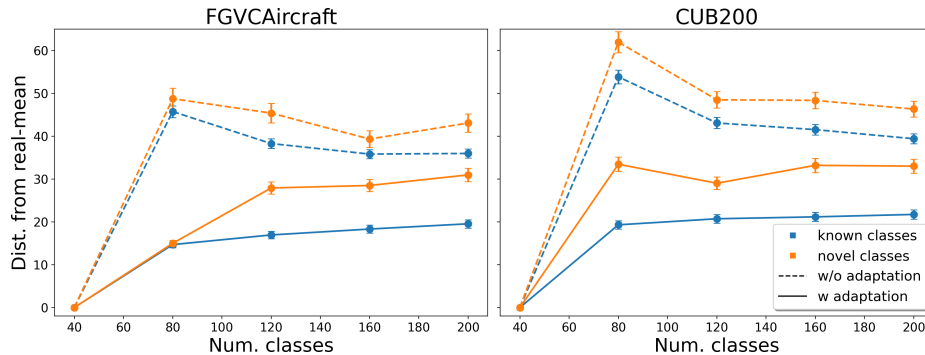


Fig. 8: Distance from centroids estimated by CAMP to ground-truth centroids throughout sequential training. Novel categories tend to be further away from estimated positions than known ones. CAMP centroid adaptation improves these estimations in both cases.

CAMP overcomes task-recency bias. To better understand the performance gap between GCD and CAMP, we check how these methods predict tasks after the last incremental step. We present a confusion matrix calculated for the StanfordCars dataset split into four tasks in Fig. 9. As a fine-tuning method, GCD faces a task-recency bias [40] for known and novel classes, as it mainly classifies test samples into categories of the last task. On the contrary, CAMP does not present such skewed results prediction only towards the last task. Our method of projected distillation, combined with category adaptation, alleviates task-recency bias, which improves CAMP’s performance.

CAMP vs. GCD under different setting conditions. We examine the impact of setting parameters (number of novel classes and fraction of labeled data) on results achieved by CAMP and GCD. We set novel to known ratios as

2:8, 4:6, 6:4, 8:2 and fraction of labeled data as 0.25, 0.5, 0.75, 1.0. We provide final all accuracy and average forgetting on StanfordCars in Fig. 10. We can observe that accuracy decreases when increasing the number of novel classes or when decreasing the amount of labeled data. Forgetting behaves similarly, which is understandable - with lower accuracy, there is less to forget. However, in each case, CAMP outperforms GCD, proving that our combination of distillation and category adaptation is robust to different scenarios.

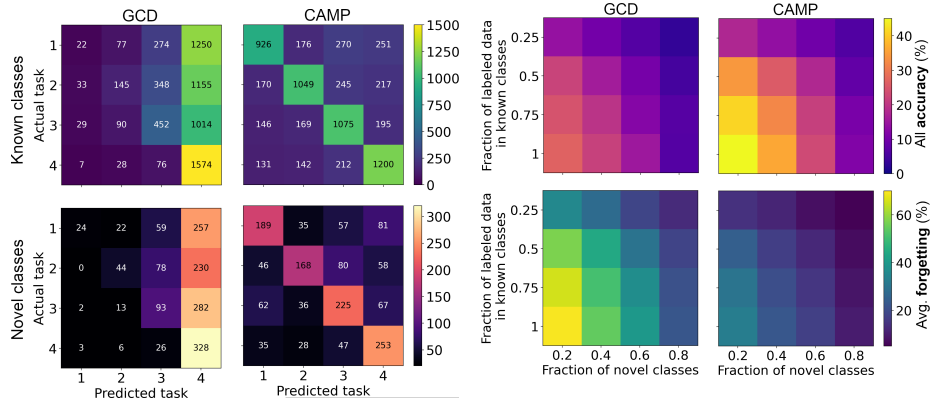


Fig. 9: Vanilla GCD method is prone to task-recency bias as its predictions are skewed towards the last one. However, CAMP overcomes this problem and keeps predictions distributed across all tasks.

Fig. 10: CAMP improves GCD in different continual learning scenarios in terms of average accuracy and forgetting. We evaluate methods for different numbers of novel classes in each task and different fractions of labeled known data on StanfordCars.

5 Summary

In this work, we study an interplay between knowledge distillation techniques and centroid adaptation in the problem of GCCD and CIL. We find out that combining knowledge distillation through a neural projector followed by centroid adaptation is crucial for achieving high plasticity and low forgetting. Based on this finding, we propose a new method, dubbed CAMP. It utilizes projected knowledge distillation to regularize the network and prevent forgetting. Moreover, CAMP utilizes a centroid adaptation method to estimate better actual positions of centroids representing past categories. CAMP uses the nearest centroid classifier to classify data samples achieving state-of-the-art results in different scenarios.

Limitations. The limitation of our work is that the centroid adaptation mechanism cannot be easily implemented for methods that do not perform centroid-based clustering to represent categories, *e.g.* [38, 45]. Moreover, our method achieves better results for known classes, which can be limiting in use cases where novel classes play the leading role.

Acknowledgments

Daniel Marczak is supported by National Centre of Science (NCN, Poland) Grant No. 2021/43/O/ST6/02482. This research was partially funded by National Science Centre, Poland, grant no: 2020/39/B/ST6/01511, 2022/45/B/ST6/02817, and 2023/51/D/ST6/02846. Bartłomiej Twardowski acknowledges the grant RYC2021-032765-I. This paper has been supported by the Horizon Europe Programme (HORIZON-CL4-2022-HUMAN-02) under the project "ELIAS: European Lighthouse of AI for Sustainability", GA no. 101120237. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2023/016613.

References

1. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: ECCV (2018)
2. Arthur, D., Vassilvitskii, S.: K-means++ the advantages of careful seeding. In: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms. pp. 1027–1035 (2007)
3. Bordes, F., Balestrieri, R., Garrido, Q., Bardes, A., Vincent, P.: Guillotine regularization: Why removing layers is needed to improve generalization in self-supervised learning. *Transactions on Machine Learning Research* (2023)
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
5. Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: Proceedings of the European conference on computer vision (ECCV). pp. 532–547 (2018)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations (2020)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR* (2021)
9. Douillard, A., Cord, M., Ollion, C., Robert, T., Valle, E.: Podnet: Pooled outputs distillation for small-tasks incremental learning. In: ECCV (2020)
10. Fini, E., da Costa, V.G.T., Alameda-Pineda, X., Ricci, E., Alahari, K., Mairal, J.: Self-supervised models are continual learners. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022. pp. 9611–9620. IEEE (2022)
11. Fini, E., Sangineto, E., Lathuilière, S., Zhong, Z., Nabi, M., Ricci, E.: A unified objective for novel class discovery. In: ICCV (2021)
12. French, R.M.: Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* **3**(4), 128–135 (1999)
13. Gomez-Villa, A., Twardowski, B., Yu, L., Bagdanov, A.D., van de Weijer, J.: Continually learning self-supervised representations with projected functional regularization (2022)
14. Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Automatically discovering and learning new visual categories with ranking statistics (2020)
15. Han, K., Rebuffi, S., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Autonovel: Automatically discovering and learning novel visual categories. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(10), 6767–6781 (2022)
16. Han, K., Vedaldi, A., Zisserman, A.: Learning to discover novel visual categories via deep transfer clustering. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8401–8409 (2019)
17. Iscen, A., Zhang, J., Lazebnik, S., Schmid, C.: Memory-efficient incremental learning through feature adaptation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16. pp. 699–715. Springer (2020)

18. Joseph, K.J., Paul, S., Aggarwal, G., Biswas, S., Rai, P., Han, K., Balasubramanian, V.N.: Novel class discovery without forgetting. In: Computer Vision - ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXIV. Lecture Notes in Computer Science, vol. 13684, pp. 570–586. Springer (2022)
19. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning (2021)
20. Kim, H., Suh, S., Kim, D., Jeong, D., Cho, H., Kim, J.: Proxy anchor-based unsupervised learning for continuous generalized category discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16688–16697 (2023)
21. Kim, S., Kim, D., Cho, M., Kwak, S.: Proxy anchor loss for deep metric learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3238–3247 (2020)
22. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* **114**(13), 3521–3526 (2017)
23. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: 2013 IEEE International Conference on Computer Vision Workshops. pp. 554–561 (2013)
24. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. . (2009)
25. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* **40**(12), 2935–2947 (2017)
26. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151* (2013)
27. Masana, M., Liu, X., Twardowski, B., Menta, M., Bagdanov, A.D., van de Weijer, J.: Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)
28. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1406–1415 (2019)
29. Petit, G., Popescu, A., Schindler, H., Picard, D., Delezoide, B.: Fetril: Feature translation for exemplar-free class-incremental learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3911–3920 (2023)
30. Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., Tesauero, G.: Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910* (2018)
31. Roy, S., Liu, M., Zhong, Z., Sebe, N., Ricci, E.: Class-incremental novel class discovery. In: European Conference on Computer Vision. pp. 317–333. Springer (2022)
32. Roy, S., Liu, M., Zhong, Z., Sebe, N., Ricci, E.: Class-incremental novel class discovery. In: Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIII. Lecture Notes in Computer Science, vol. 13693, pp. 317–333. Springer (2022)
33. Rypeść, G., Cygert, S., Khan, V., Trzciński, T., Zieliński, B., Twardowski, B.: Divide and not forget: Ensemble of selectively trained experts in continual learning. *arXiv preprint arXiv:2401.10191* (2024)

34. Sariyildiz, M.B., Kalantidis, Y., Alahari, K., Larlus, D.: No reason for no supervision: Improved generalization in supervised models. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023. OpenReview.net (2023)
35. Sariyildiz, M.B., Kalantidis, Y., Alahari, K., Larlus, D.: No reason for no supervision: Improved generalization in supervised models. In: ICLR 2023-International Conference on Learning Representations. pp. 1–26 (2023)
36. Vaze, S., Han, K., Vedaldi, A., Zisserman, A.: Generalized category discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7492–7501 (2022)
37. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
38. Wen, X., Zhao, B., Qi, X.: Parametric classification for generalized category discovery: A baseline study. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16590–16600 (2023)
39. Wu, Y., Chi, Z., Wang, Y., Feng, S.: Metagcd: Learning to continually learn in generalized category discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1655–1665 (2023)
40. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Fu, Y.: Large scale incremental learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 374–382 (2019)
41. Yan, S., Xie, J., He, X.: Der: Dynamically expandable representation for class incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3014–3023 (2021)
42. Yu, L., Twardowski, B., Liu, X., Herranz, L., Wang, K., Cheng, Y., Jui, S., Weijer, J.v.d.: Semantic drift compensation for class-incremental learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6982–6991 (2020)
43. Zhang, S., Khan, S., Shen, Z., Naseer, M., Chen, G., Khan, F.S.: Promptcal: Contrastive affinity learning via auxiliary prompts for generalized novel category discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3479–3488 (2023)
44. Zhang, X., Jiang, J., Feng, Y., Wu, Z.F., Zhao, X., Wan, H., Tang, M., Jin, R., Gao, Y.: Grow and merge: A unified framework for continuous categories discovery. *Advances in Neural Information Processing Systems* **35**, 27455–27468 (2022)
45. Zhao, B., Mac Aodha, O.: Incremental generalized category discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19137–19147 (2023)
46. Zhao, B., Wen, X., Han, K.: Learning semi-supervised gaussian mixture models for generalized category discovery. *arXiv preprint arXiv:2305.06144* (2023)
47. Zhong, Z., Fini, E., Roy, S., Luo, Z., Ricci, E., Sebe, N.: Neighborhood contrastive learning for novel class discovery. In: CVPR (2021)
48. Zhong, Z., Zhu, L., Luo, Z., Li, S., Yang, Y., Sebe, N.: Openmix: Reviving known knowledge for discovering novel visual categories in an open world (2020)
49. Zhou, D.W., Wang, F.Y., Ye, H.J., Zhan, D.C.: Pycil: a python toolbox for class-incremental learning. *SCIENCE CHINA Information Sciences* **66**(9), 197101– (2023)
50. Zhu, F., Cheng, Z., Zhang, X.y., Liu, C.l.: Class-incremental learning via dual augmentation. *Advances in Neural Information Processing Systems* **34** (2021)
51. Zhu, F., Zhang, X.Y., Wang, C., Yin, F., Liu, C.L.: Prototype augmentation and self-supervision for incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5871–5880 (2021)

52. Zhu, K., Zhai, W., Cao, Y., Luo, J., Zha, Z.J.: Self-sustaining representation expansion for non-exemplar class-incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9296–9305 (2022)