

Class-incremental Novel Class Discovery

Subhankar Roy^{1,2}, Mingxuan Liu¹, Zhun Zhong¹,
Nicu Sebe¹, and Elisa Ricci^{1,2}

¹ University of Trento, Trento, Italy

² Fondazione Bruno Kessler, Trento, Italy

{subhankar.roy, mingxuan.liu, zhun.zhong, niculae.sebe, e.ricci}@unitn.it

Abstract. We study the new task of class-incremental Novel Class Discovery (class-iNCD), which refers to the problem of discovering novel categories in an unlabelled data set by leveraging a pre-trained model that has been trained on a labelled data set containing disjoint yet related categories. Apart from discovering novel classes, we also aim at preserving the ability of the model to recognize previously seen base categories. Inspired by rehearsal-based incremental learning methods, in this paper we propose a novel approach for class-iNCD which prevents forgetting of past information about the base classes by jointly exploiting base class feature prototypes and feature-level knowledge distillation. We also propose a self-training clustering strategy that simultaneously clusters novel categories and trains a joint classifier for both the base and novel classes. This makes our method able to operate in a class-incremental setting. Our experiments, conducted on three common benchmarks, demonstrate that our method significantly outperforms state-of-the-art approaches. Code is available at <https://github.com/OatmealLiu/class-iNCD>.

Keywords: Novel Class Discovery, Class-Incremental Learning

1 Introduction

Humans are bestowed with the excellent cognitive skills to learn continually over their lifetime [12], and in most cases without the need of explicit supervision [1]. Thus, it has been a long-standing goal of the machine learning research community to build Artificial Intelligence (AI) systems that can mimic this human-level performance. In an attempt to realize this, much effort has been dedicated to learn deep learning models from large reservoirs of both labelled [24,17,10] and unlabelled data [3,4]. Aside from being effective learners, by imitating human learning mechanisms, neural networks should also be flexible to absorb novel concepts (or *classes*) after having learned some patterns with the past data. The task of automatically discovering novel (or *new*) classes in an unsupervised fashion while leveraging some previously learned knowledge is referred to as *novel class discovery* (NCD) [15,16,37,38,11] (see Fig. 1(a)). NCD has gained significant attention in the recent times due to its practicality of efficiently learning novel classes without relying on large quantities of unlabelled data [15].

The first two authors contributed equally. Corresponding author: Zhun Zhong

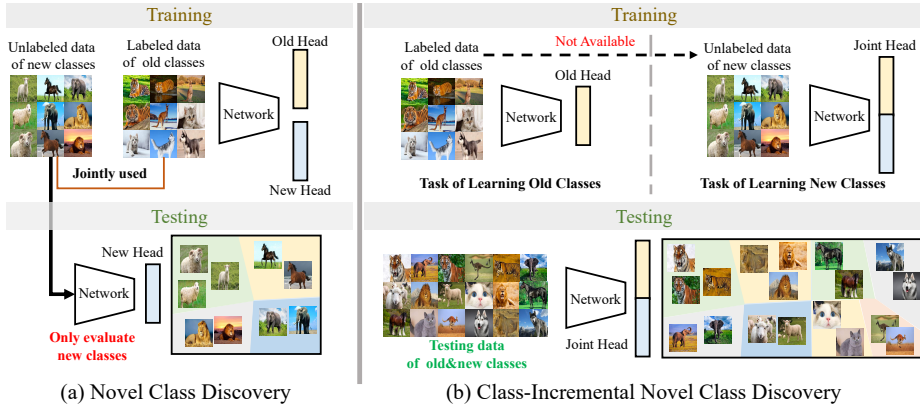


Fig. 1. Comparison between the settings (a) Novel Class Discovery (NCD) which solely concerns the performance of novel classes, and (b) the proposed class-incremental NCD (class-iNCD) measures performance of all the classes seen so far with a single classifier.

Most of the proposed NCD solutions rely on stage-wise [19,20,16] or joint [15,37,11] learning schemes on the labelled and the unlabelled data, with the assumption that structures discovered on the labelled images could be leveraged as a proxy supervision on the unlabelled images. It has been shown that NCD benefits more when the model is trained jointly on the labelled data while using a clustering objective on the unlabelled data [15,38,37,11]. However, access to the labelled data after the pre-training stage can not always be guaranteed in real-world applications due to privacy or storage issues. This calls for a more pragmatic NCD setting where the labelled images would be discarded and only the pre-trained model could be transferred for learning the novel classes. Being meaningful, such *source-free* model adaptation has been explored in the related areas of domain adaptation [28,36]. Although seems more practical, such a training scheme would gradually cause the network to erase all the previously learned information about the old (or *base*) classes. This drop in the base class performance when the labelled data set becomes unavailable is primarily attributed to the phenomenon of *catastrophic forgetting* [8] in neural networks. In most of the aforementioned NCD methods the performance on the novel classes are only deemed important, without any consideration for preserving the performance on the base classes. We believe that such a setting is of little practical significance in the real world because the adapted model becomes unusable on the base classes and retraining is infeasible.

Given the inherent drawbacks of the existing NCD setting, we argue that an ideal NCD method should aim to learn novel classes without the explicit presence of the labelled data and at the same time preserve the performance on the base classes. This new setting is referred to as *task-incremental* NCD (iNCD), and indeed has been very recently studied in [29]. In details, ResTune [29] uses knowledge distillation [27] on the network logits to prevent forgetting on the

base classes and a clustering objective [33] with task specific network weights for the novel classes. As opposed to the ResTune [29], which facilitates iNCD by solely improving the ability of the network to learn novel classes, we additionally improve the incremental learning aspect in iNCD as well. Specifically, inspired by the rehearsal-based incremental learning methods [2,7,31] which are known to be effective, we propose to store the base class feature *prototypes* from the previous task as exemplars, instead of raw images. Features derived from the stored prototypes are then *replayed* to prevent forgetting old information on the base classes in addition to feature-level knowledge distillation. On the other hand, to facilitate learning of novel classes, we dedicate a task specific classifier that is optimized with robust rank statistics [15]. Disadvantageously, the introduction of task specific classifier leads to the dependence on the task-id of an input sample during inference. To overcome reliance on task-id, we propose to maintain a joint classifier for both the base and novel classes, which is trained with the pseudo-labels generated by the task specific one. We call this setting as *class-incremental* NCD (class-iNCD) as it does not allow the task-id information to be used during inference. The high level overview of the new class-iNCD setting is shown in the Fig. 1(b). As our proposed method amalgamates **Feature Replay** and **Distillation** with **Self-Training** we name it FRoST. In summary, the contributions of this work are three-fold:

- We propose a novel framework, FRoST, that can tackle the newly introduced and relevant task of class-incremental novel class discovery (class-iNCD).
- Our FRoST is equipped with prototypes for feature-replay and employs feature-level knowledge distillation to prevent forgetting. Moreover, it uses pseudo-labels from the task specific head to efficiently learn novel classes without interference, enabling us to achieve a task-agnostic classifier.
- We run extensive experiments on three common benchmarks to prove the effectiveness of our method. FRoST also obtains state-of-the-art performance when compared with the existing baselines. Additionally, we run experiments on a sequence of tasks of unlabelled sets and verify its generality.

2 Related Works

Novel Class Discovery (NCD) deals with the task of learning to discover new semantic classes in an unlabelled data set by utilizing the knowledge acquired from another labelled data set [16,35]. It is assumed that the classes in the labelled and unlabelled set are disjoint. So far, several NCD methods have been proposed and they can be broadly classified into two broad sub-categories. The first category of NCD methods use a stage-wise training scheme where the model is first pre-trained on the labelled set, followed by fine-tuning on the unlabelled data using an unsupervised clustering loss [19,20,16,29]. Barring [29], none of the above methods consider to tackle the forgetting issue, and as a result the model loses the ability to classify the base classes. The second category comprise of NCD methods that assume both the labelled and unlabelled data are available simultaneously, which are then trained jointly [15,37,38,21,11]. As demonstrated

in [29], the NCD methods which rely on joint training always outperform the stage-wise NCD methods. However, the latter family of NCD methods rely on the availability of labelled data, which is often not permitted due to privacy reasons. This makes stage-wise training scheme favourable to tackle class-iNCD, but it lacks the capability to prevent forgetting. Similar to the ResTune [29], we also build our framework that can be trained in a stage-wise manner and also be able to maintain performance on the base classes. Different from the ResTune we use the predictions of the novel class classifier as pseudo-labels (PL) to train a single joint classifier that can classify both base and novel classes.

Incremental Learning (IL) is a learning paradigm where a model is trained on a sequence of tasks such that data from only the current task is available for training, while the model is evaluated on all the observed tasks. The IL methods are designed so as to prevent catastrophic forgetting [13] of the model on the old tasks and at the same time flexible enough to learn on new tasks [6]. Most early IL methods addressed the *task-incremental learning* setting (task-IL), where the model has access to a task-id for choosing the task-specific classifier during the testing phase. Given the practical limitations of knowing the task-id during inference, more recent IL methods have started to address the *class-incremental learning* (class-IL) setting, where the task-id is not available during inference. This makes class-IL setting practical and at the same time more challenging than the task-IL setting. Our FRoST also operates in the class-IL setting, which we call as class-iNCD. Existing IL methods can be sub-divided into three broad categories: *regularization-based* methods [22,34,27,9], *exemplar-based* methods [31,5,2,7] and methods focused on *task-recency bias* problem [32]. We refer the readers to the survey in [30] for an exhaustive list of class-IL methods. In our FRoST we propose to use a combination of knowledge distillation [27] at intermediate feature-level and storage of base class feature prototypes as exemplars to prevent forgetting in feature extractor and classifier, respectively. We discuss later in Sec. 3 why this choice is suitable for the class-iNCD setting.

3 Method

In this section we describe our FRoST for the task of class-iNCD. Before delving into the detail we lay down some preliminaries related to our method.

Problem Definition and Notation. In the setting of class-incremental novel class discovery (class-iNCD) we are initially given $n^{[L]}$ instances of a labelled data set $\mathcal{D}^{[L]} = \{(\mathbf{x}_i^{[L]}, \mathbf{y}_i^{[L]})\}_{i=1}^{n^{[L]}}$ belonging to the supervised task $\mathcal{T}^{[L]}$, where $\mathbf{x}^{[L]} \in \mathcal{X}^{[L]}$ represents the input images and $\mathbf{y}^{[L]} \in \mathcal{Y}^{[L]}$ as $|\mathcal{C}^{[L]}|$ -dimensional one-hot labels. Once standard supervised training is finished on the task $\mathcal{T}^{[L]}$, the data set $\mathcal{D}^{[L]}$ is discarded and we are presented with $n^{[U]}$ instances from a new task $\mathcal{T}^{[U]}$. The task $\mathcal{T}^{[U]}$ has an unlabelled data set $\mathcal{D}^{[U]} = \{\mathbf{x}_j^{[U]}\}_{j=i}^{n^{[U]}}$ where $\mathbf{x}^{[U]} \in \mathcal{X}^{[U]}$ are the unlabelled images containing $\mathcal{C}^{[U]}$ classes. As in any NCD setting [16], it is assumed that the labels in $\mathcal{Y}^{[L]}$ and $\mathcal{Y}^{[U]}$ are disjoint, *i.e.*, $\mathcal{Y}^{[L]} \cap \mathcal{Y}^{[U]} = \emptyset$. The goal of class-iNCD is to cluster the images in $\mathcal{D}^{[U]}$ by just leveraging the learnt information contained in the mapping function $f^{[L]}: \mathcal{X}^{[L]} \rightarrow \mathcal{Y}^{[L]}$, while

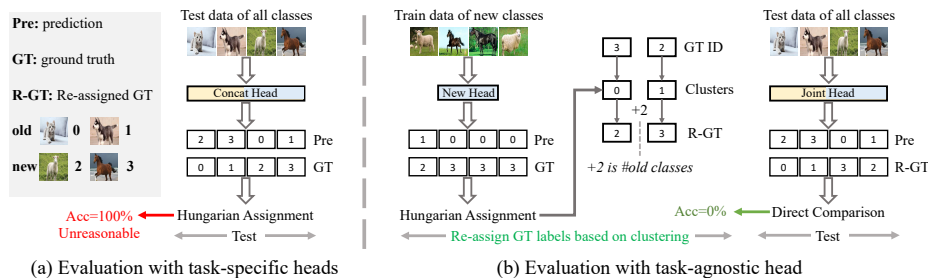


Fig. 2. Evaluation protocol comparison (a) evaluation with task-specific heads in iNCD [29] and (b) evaluation with task-agnostic head in our class-iNCD.

still behaving well on the previous task $\mathcal{T}^{[L]}$. In other words, we are interested in learning a single mapping function $f: \mathcal{X} \rightarrow \mathcal{Y}^{[L]} \cup \mathcal{Y}^{[U]}$ that can be used to infer the label of any test image $\mathbf{x} \in \{\mathcal{X}^{[L]} \cup \mathcal{X}^{[U]}\}$. This is in sharp contrast to the existing NCD methods where the performance on $\mathcal{T}^{[L]}$ is not of interest.

Evaluation Protocol. In the NCD methods [11,37], task-specific heads are trained for old and new classes³, respectively. This poses a limitation, as they can only operate in task-specific NCD setting. To address this problem, ResTune [29] uses the concatenation of old and new heads during inference. The class-incremental performance is estimated with the Hungarian Assignment (HA) [25] by regarding this problem as a clustering task. However, this evaluation protocol is indeed improper in class-iNCD, since it does not explicitly distinguish the old and new classes. As shown in the Fig. 2(a), the classifier recognizes the samples of old classes as novel classes (and vice versa), and yet the accuracy obtained by HA is still 100%, making the evaluation in [29] unfair.

In this work, we learn a task-agnostic head (or *joint* head) and propose a new evaluation protocol for class-iNCD (see Fig. 2(b)). In details, we first use the *new head* to estimate the predictions of unlabeled data from the new classes. We utilize the HA [25] to re-assign ground-truth IDs based on the predictions and ground-truth labels for the new classes only. The joint (task-agnostic) classifier is used to evaluate the new classes test samples by directly comparing the predictions with these re-assigned ground-truth labels. Whereas for the old classes test data, we evaluate using the old classes ground truth. As shown in Fig. 2(b), our evaluation protocol explicitly distinguishes the old and new classes. As evident, our evaluation is more reasonable than [29] and penalizes the metric when the new classes are classified as one of the old classes, which is an ideal behaviour.

Overall Framework. Being in the incremental learning setting, our proposed FRoST (see Fig. 3) operates in two stages. In the first stage we learn the mapping function $f^{[L]}: \mathcal{X}^{[L]} \rightarrow \mathcal{Y}^{[L]}$ in a supervised manner on the labelled data set $\mathcal{D}^{[L]}$ that can recognize samples belonging to the first $\mathcal{C}^{[L]}$ categories. We model the function $f^{[L]}$ with a neural network that is further composed of two sub-networks: feature extractor $g(\cdot)$ and a linear classifier $h^{[L]}(\cdot)$ that outputs $\mathcal{C}^{[L]}$ logits, such

³ When referring to classes, we regard old & base; and, new & novel interchangeably.

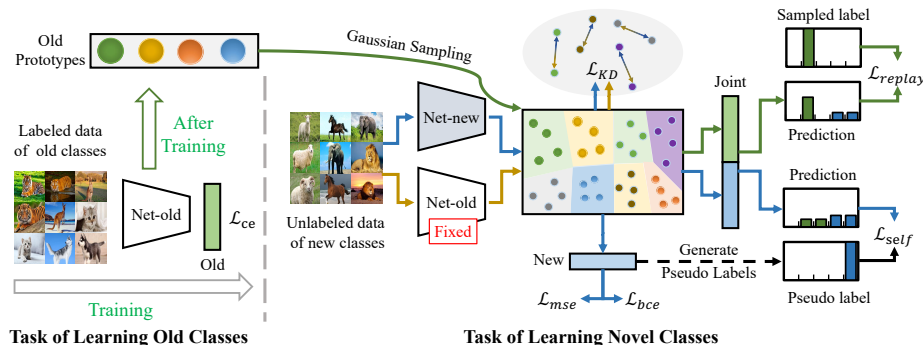


Fig. 3. An overview of the proposed FroST. **Left:** a base model is learned supervisedly (\mathcal{L}_{ce}) on the old classes. Old class-prototypes and variances are stored. **Right:** the new classes are learned with a clustering objective (\mathcal{L}_{bce}). Forgetting on old classes is prevented by using feature-distillation (\mathcal{L}_{KD}) and feature-replay (\mathcal{L}_{replay}) with the class-prototypes. A joint classifier is learned by self-training (\mathcal{L}_{self}) with pseudo-labels.

that $f^{[L]} = h^{[L]} \circ g$. The feature extractor g and classifier $h^{[L]}$ are parameterized by θ_g and $\theta_{h^{[L]}}$, respectively. Before we move to the second stage, we compute per-class intermediate feature prototypes μ_c from the intermediate features $\mathbf{z}^{[L]} = g(\mathbf{x}^{[L]})$, belonging to each class c . Additionally, we also compute and store the variance of the features of class c as \mathbf{v}_c^2 .

In the second stage, the $\mathcal{D}^{[L]}$ is discarded and the novel classes are learned on $\mathcal{D}^{[U]}$ by reusing the transferred network weights $f^{[L]}$. Since our goal is to learn an unique classifier that can accommodate $\mathcal{C}^{[A]} = \mathcal{C}^{[L]} + \mathcal{C}^{[U]}$ classes, we extend the classifier $h^{[L]}$ to $h^{[A]}$ in order to incorporate the $\mathcal{C}^{[U]}$ novel classes. Besides $h^{[A]}$, we instantiate a new task-specific classifier $h^{[U]}$ for $\mathcal{T}^{[U]}$ that is trained on $\mathcal{D}^{[U]}$ to exclusively classify the novel classes. The classifiers $h^{[A]}$ and $h^{[U]}$ are parameterized by $\theta_{h^{[U]}}$ and $\theta_{h^{[A]}}$, respectively. In details, the network $f^{[U]} = h^{[U]} \circ g$ is trained using the clustering objective in [15] that leverages previously learned information to provide supervision using the robust rank statistics. With the goal of learning a joint classifier $h^{[A]}$, we obtain pseudo-label for $\mathbf{x}^{[U]}$ from $h^{[U]}$ and distill it to the newly extended part of $h^{[A]}$, which handles the novel classes. On the other hand, to mitigate forgetting on the base classes of $\mathcal{T}^{[L]}$ we employ two strategies: *feature-level knowledge distillation* [18,27] on g that ensures the feature encoding for the old task $\mathcal{T}^{[L]}$ does not drift too far while learning on $\mathcal{T}^{[U]}$; and *generative feature-replay* drawn from a Gaussian distribution $\mathcal{N}(\mu_c, \mathbf{v}_c^2)$ is used to preserve performance of the top part of the $h^{[A]}$, which is responsible for classifying the base classes. During inference the classifier $h^{[A]}$ is used.

3.1 Preliminaries

Supervised Training. In the first stage of the class-iNCD task we have at disposal the labelled images from $\mathcal{D}^{[L]}$. This stage consists in learning a supervised

model $f^{[L]}$ that can classify the base classes drawn from the task $\mathcal{T}^{[L]}$. We aim to learn the parameters $(\{\theta_g, \theta_{h^{[L]}}\})$ of the model $f^{[L]} = h^{[L]} \circ g$ by using a supervised *cross-entropy* loss:

$$\mathcal{L}_{ce} = -\mathbb{E}_{p(\mathbf{x}^{[L]}, \mathbf{y}^{[L]})} \frac{1}{C^{[L]}} \sum_{k=1}^{C^{[L]}} y_k^{[L]} \log \sigma_k(h^{[L]}(g(\mathbf{x}^{[L]}))), \quad (1)$$

where $\sigma_k(\mathbf{1}) = \exp(l_k) / \sum_j \exp(l_j)$ represents the likelihood corresponding to the k^{th} output from the model and $C^{[L]}$ is the number of classes in the task $\mathcal{T}^{[L]}$.

Knowledge Distillation to Prevent Forgetting. Having learned an optimal model on a given task, the main challenge in IL is to learn new tasks without forgetting the past information. A very popular regularization-based approach to overcome forgetting on previously learned tasks is by using *knowledge distillation* (KD) [18]. Concretely, based on KD, *Learning without Forgetting* (LwF) [27] is an effective method commonly used in IL. It consists in penalizing the network if the representation of data from previous tasks drifts too far while learning on a new task. Assuming a simplified task-IL learning scenario containing just two tasks: $\mathcal{T}^{[\text{old}]}$ and $\mathcal{T}^{[\text{new}]}$, where a model $f^{[\text{old}]} = h^{[\text{old}]} \circ g$ has already been trained using the objective in Eq. (1) and a new task $\mathcal{T}^{[\text{new}]}$ has been presented to the learning algorithm. The goal of LwF is to prevent forgetting on $\mathcal{T}^{[\text{old}]}$ while learning on $\mathcal{T}^{[\text{new}]}$. LwF keeps a copy of the old model $f^{[\text{old}]} = h^{[\text{old}]} \circ g^{[\text{old}]}$ and simultaneously creates a new instance $f^{[\text{new}]} = h^{[\text{new}]} \circ g^{[\text{new}]}$ (with $g^{[\text{new}]} = g^{[\text{old}]}$) for learning on $\mathcal{T}^{[\text{new}]}$. The $f^{[\text{new}]}$ differs from $f^{[\text{old}]}$ in the final classification head where the task-specific classifier $h^{[\text{new}]}$ exclusively handles the class assignment for the new classes in $\mathcal{T}^{[\text{new}]}$. Given a sample $\mathbf{x}^{[\text{new}]}$ from the new task $\mathcal{T}^{[\text{new}]}$, LwF aims to match the pre-recorded logits $\mathbf{a}^{[\text{old}]} = h^{[\text{old}]}(g^{[\text{old}]}(\mathbf{x}^{[\text{new}]}))$ from the frozen $f^{[\text{old}]}$ with the old task logits $\hat{\mathbf{a}}^{[\text{old}]} = h^{[\text{old}]}(g^{[\text{new}]}(\mathbf{x}^{[\text{new}]}))$. Essentially, this prevents $g^{[\text{new}]}$ to produce feature encoding that is too different from that of $g^{[\text{old}]}$, since the success for the old task heavily depends on it. Formally, the LwF loss at logits-level is given as:

$$\mathcal{L}_{\text{KD}}^{\text{logits}} = -\mathbb{E}_{p(\mathbf{x}^{[\text{new}]})} \frac{1}{K^{[\text{old}]}} \sum_{k=1}^{K^{[\text{old}]}} \pi_k(h^{[\text{old}]}(g^{[\text{old}]}(\mathbf{x}^{[\text{new}]})) \log \pi_k(h^{[\text{old}]}(g^{[\text{new}]}(\mathbf{x}^{[\text{new}]})))), \quad (2)$$

where $\pi_k(\mathbf{a}) = \exp(a_k/\tau) / \sum_j \exp(a_j/\tau)$ is the temperature controlled likelihood of the model with τ being the temperature. The parameters $(\{\theta_{g^{[\text{new}]}}\}, \theta_{h^{[\text{old}]}}\})$ corresponding to $g^{[\text{new}]}$ and $h^{[\text{old}]}$ are updated with Eq. (2). However, the need of having separate task-specific classifiers in the LwF approach limits the applicability of such models to the task-IL setting, as in ResTune [29]. While LwF can ideally be extended to the class-IL consisting of a single classifier, it would require pre-allocation of all the logits during the first task. The assumption of knowing a priori the cardinality of the tasks and their constituent classes is impractical. Thus, we build on top LwF and adapt it to the class-iNCD.

3.2 Class-incremental Novel Class Discovery

We are interested in learning a model that can incrementally cluster unlabelled images into a set of novel classes, after the model has been trained on a labelled set of images. Besides good performance on the novel classes we also desire to preserve the performance on the previously seen classes, without having access to or storing images from the previous tasks. Most importantly, at any point of time during the training sessions, we maintain a single classification head for all the classes seen so far. To address the challenging task of class-iNCD we propose to tackle it from two different axes. The first axis is concerned with learning discriminative features on the unlabelled data set $\mathcal{D}^{[u]}$ by using a clustering objective. Although the model gets better at classifying the novel classes, its performance gradually deteriorates on the base classes due to forgetting [8]. To overcome this issue, the second axis deals with preventing forgetting on all base classes by using the images only from the new task, combined with a feature-replay strategy. We elaborate them below.

Self-training for Novel Class Discovery. When presented with an unlabelled data set $\mathcal{D}^{[u]}$, the discovery step in class-iNCD involves learning the weights of the network $f^{[u]} = h^{[u]} \circ g$. While the newly initialized classifier $h^{[u]}$ yet lacks the capability to classify images into novel categories, the feature extractor g on the other hand has already been trained on a related labelled data set $\mathcal{D}^{[L]}$ and has a notion of what constitutes a semantic concept in an image. Adopting this ideology from the NCD method AutoNovel [15], the pairwise similarity between a pair of unlabelled images $(\mathbf{x}_i^{[u]}, \mathbf{x}_j^{[u]})$ is inferred and provided as a weak form of supervision in the discovery step. The feature descriptors $\mathbf{z}_i^{[u]} = g(\mathbf{x}_i^{[u]})$ and $\mathbf{z}_j^{[u]} = g(\mathbf{x}_j^{[u]})$ corresponding to the pair $(\mathbf{x}_i^{[u]}, \mathbf{x}_j^{[u]})$ are then compared using the robust rank statistics. If the top- k ranked dimensions of the feature descriptor pair $(\mathbf{z}_i^{[u]}, \mathbf{z}_j^{[u]})$ are found the same then $(\mathbf{x}_i^{[u]}, \mathbf{x}_j^{[u]})$ can be considered to belong to the same class. The pairwise pseudo-label is formulated as:

$$\tilde{y}_{ij}^{[u]} = \mathbb{1}\{\text{top}_k(\mathbf{z}_i^{[u]}) = \text{top}_k(\mathbf{z}_j^{[u]})\}, \quad (3)$$

where $\text{top}_k : \mathbf{z}^{[u]} \rightarrow \mathcal{S}\{1, \dots, k\} \subset \mathcal{P}\{1, \dots, |\mathbf{z}^{[u]}|\}$ denotes the subset of top- k most activated feature indices in $\mathbf{z}^{[u]}$. This pairwise pseudo-label is then used to train the classifier $h^{[u]}$ for the novel classes. In detail, the dot-product of the classifier’s predictions $p_{ij} = \sigma(\langle h^{[u]}(g(\mathbf{x}_i^{[u]})), h^{[u]}(g(\mathbf{x}_j^{[u]})) \rangle)$ can be interpreted as a similarity between $\mathbf{x}_i^{[u]}$ and $\mathbf{x}_j^{[u]}$, where $\sigma(\cdot)$ is a logistic function. Thus, the pairwise pseudo-label $\tilde{y}_{ij}^{[u]}$ computed in Eq. (3) is used to enforce this association between $\mathbf{x}_i^{[u]}$ and $\mathbf{x}_j^{[u]}$. The parameters $(\{\theta_g, \theta_{h^{[u]}}\})$ are trained with a *binary cross-entropy* loss as:

$$\mathcal{L}_{\text{bce}} = -\mathbb{E}_{p(\mathbf{z}^{[u]})} \tilde{y}_{ij}^{[u]} \log(p_{ij}) + (1 - \tilde{y}_{ij}^{[u]}) \log(1 - p_{ij}). \quad (4)$$

While the objective in Eq. (4) learns a classifier for the novel classes, such training scheme makes the inference step dependent on task-id like ResTune. In

order to make our model suitable for class-iNCD we resort to self-training with the help of pseudo-labels that are computed from $f^{[U]}$ to train the joint classifier $h^{[A]}$. In details, given the goal of learning the model $f^{[A]} = h^{[A]} \circ g$, we use $h^{[U]}$ to compute the pseudo-label $\hat{y}^{[U]}$ for an unlabelled image $\mathbf{x}^{[U]}$. The $\hat{y}^{[U]}$ is then used to supervise the training of $h^{[A]}$. The self-training loss is described as:

$$\mathcal{L}_{\text{self}} = -\mathbb{E}_{(\mathbf{x}^{[U]}, \hat{\mathbf{y}}^{[U]})} \frac{1}{|C^{[A]}|} \sum_{k=1}^{|C^{[A]}|} \hat{y}_k^{[U]} \log \sigma_k(h^{[A]}(g(\mathbf{x}^{[U]}))), \quad (5)$$

where

$$\hat{y}^{[U]} = C^{[L]} + \arg \max_{k \in C^{[U]}} h^{[U]}(g(\mathbf{x}^{[U]})). \quad (6)$$

Since the pairwise pseudo-labels in Eq. (3) can be noisy, it can lead to a poorly trained $h^{[U]}$. As a consequence, the noisy pseudo-labels $\hat{y}^{[U]}$ from $h^{[U]}$ can have an adverse impact on the training of the joint classifier $h^{[A]}$. To minimize the cascading error propagation we also enforce consistency between two correlated views for an unlabelled image $\mathbf{x}^{[U]}$. Specifically, using stochastic data augmentation on $\mathbf{x}^{[U]}$ we generate a correlated view $\bar{\mathbf{x}}^{[U]}$ and optimize $(\{\theta_g, \theta_{h^{[U]}}\})$ with a *mean-squared error* loss as:

$$\mathcal{L}_{\text{mse}} = \mathbb{E}_{p(\mathbf{x}^{[U]}, \bar{\mathbf{x}}^{[U]})} \frac{1}{|C^{[U]}|} \sum_{k=1}^{|C^{[U]}|} \left(\sigma_k(h^{[U]}(g(\mathbf{x}^{[U]}))) - \sigma_k(h^{[U]}(g(\bar{\mathbf{x}}^{[U]}))) \right)^2. \quad (7)$$

Finally, the overall loss for discovering novel classes and having a single classifier for all the classes seen so far can be written as:

$$\mathcal{L}_{\text{novel}} = \mathcal{L}_{\text{bce}} + \omega_{\text{self}}(t) \mathcal{L}_{\text{self}} + \omega_{\text{mse}}(t) \mathcal{L}_{\text{mse}}, \quad (8)$$

where $\omega_{\text{self}}(t)$ and $\omega_{\text{mse}}(t)$ are ramp-up functions to ensure stability in learning. **Feature Replay and Distillation for Class-incremental Learning.** While the proposed self-training assists the model $f^{[A]}$ in discovering the novel classes, it simultaneously loses the ability to predict the old classes in $\mathcal{T}^{[U]}$. To mitigate the forgetting we propose feature replay and feature distillation. To recall, at the end supervised training on $\mathcal{T}^{[L]}$ and before discarding $\mathcal{D}^{[L]}$ we compute the class prototype $\boldsymbol{\mu}_c^{[L]}$ and variance $\mathbf{v}_c^{[L]2}$ for each base class as:

$$\boldsymbol{\mu}_c^{[L]} = \frac{1}{n_c^{[L]}} \sum_{i=1}^{n_c^{[L]}} g(\mathbf{x}_i^{[L]}), \quad \mathbf{v}_c^{[L]2} = \frac{1}{n_c^{[L]}} \sum_{i=1}^{n_c^{[L]}} (g(\mathbf{x}_i^{[L]}) - \boldsymbol{\mu}_c^{[L]})^2, \quad (9)$$

where $n_c^{[L]}$ denotes the number of samples belonging to class c in $\mathcal{D}^{[L]}$. While learning on the new task $\mathcal{T}^{[U]}$; the weights of the joint classifier $h^{[A]}$, corresponding to the base classes $C^{[L]}$, are trained by replaying features from the class-specific Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_c^{[L]}, \mathbf{v}_c^{[L]2})$ of $\mathcal{T}^{[L]}$. The feature-replay loss is given as:

$$\mathcal{L}_{\text{replay}} = -\mathbb{E}_{c \sim C^{[L]}} \mathbb{E}_{(\mathbf{z}^{[L]}, \mathbf{y}_c^{[L]}) \sim \mathcal{N}(\boldsymbol{\mu}_c, \mathbf{v}_c^2)} \sum_{k=1}^{|C^{[A]}|} y_{kc}^{[L]} \log \sigma_k(h^{[A]}(\mathbf{z}^{[L]})). \quad (10)$$

As the feature extractor g also gets updated during the optimization of Eq. (8), this will make the prototypes outdated. To keep the feature replay useful we add an extra regularization on g with the help of feature distillation, which is given as:

$$\mathcal{L}_{\text{KD}}^{\text{feat}} = -\mathbb{E}_{p(\mathbf{x}^{[u]})} \left\| g^{[L]}(\mathbf{x}^{[u]}) - g(\mathbf{x}^{[u]}) \right\|_2, \quad (11)$$

where $g^{[L]}$ is the feature extractor from the previous task and is kept frozen.

Conventionally, in supervised class-IL or task-IL, regularization with the LwF loss in the logits space while learning supervisedly on a new task is usually effective in preventing forgetting. Contrarily in class-iNCD, as the novel classes need to be learned without explicit supervision, it makes the optimization of NCD part interfere with that of forgetting. This motivates us to decouple the objective for *not-forgetting* into $\mathcal{L}_{\text{replay}}$ and $\mathcal{L}_{\text{KD}}^{\text{feat}}$. We show later in Sec. 4 with adequate experiments the disadvantages of using LwF on the logits of the network. The overall objective for not-forgetting past information is given as:

$$\mathcal{L}_{\text{past}} = \mathcal{L}_{\text{replay}} + \lambda \mathcal{L}_{\text{KD}}^{\text{feat}}, \quad (12)$$

where λ is used to weight the feature distillation loss.

Overall Training. Finally, our FRoST is optimized with the following objective:

$$\mathcal{L}_{\text{FRoST}} = \mathcal{L}_{\text{novel}} + \mathcal{L}_{\text{past}}. \quad (13)$$

4 Experiments

4.1 Experimental Setup

Datasets. We have used three data sets to conduct experiments for class-iNCD: CIFAR-10 [23], CIFAR-100 [23] and Tiny-ImageNet [26]. We split the data sets into the old and new classes following the existing NCD and iNCD works [15,37,29]. The splits are reported in the supplementary material.

Evaluation metrics. We used our new evaluation protocol (Sec. 3) to evaluate the performance on the test data for all the classes. We report three classification accuracies, denoted as **Old**, **New** and **All**. They represent the accuracy obtained from the joint classifier head on the samples of the old, new and old+new classes, respectively. Refer to the supplementary material for details.


Implementation details We used ResNet-18 [17] as the backbone in all the experiments. We have adopted most of the hyperparameters from AutoNovel [15]. We introduce only one additional hyperparameter λ , which is set to 10. Rest of the implementation details can be found in the supplementary material.

4.2 Ablation Studies

Effectiveness of Feature Replay and Distillation. In Tab. 1 we measure the impact of the components proposed for not forgetting: feature distillation (FD),

Table 1. Ablation study of the proposed feature distillation (FD), feature replay (FR) and self-training (ST) that form our FRoST for the class-iNCD.

Methods	CIFAR-10			CIFAR-100			Tiny-ImageNet			Average		
	Old	New	All	Old	New	All	Old	New	All	Old	New	All
FRoST (Ours)	77.4	49.5	63.5	62.5	45.8	59.2	54.4	33.9	52.4	64.8	43.1	58.3
w/o FD & FR	0.0	36.4	18.2	0.0	33.1	6.6	0.0	37.2	3.7	0.0	35.6	9.5
w/o FD	0.0	39.4	19.7	0.0	33.1	6.6	0.0	34.3	3.4	0.0	35.6	9.9
w/o FR	0.0	73.3	36.6	0.0	57.8	11.6	0.0	40.9	4.1	0.0	57.3	17.4
w/o ST	91.7	0.0	45.8	69.2	0.0	55.4	57.5	0.0	51.7	72.8	0.0	51.0
w/o FD & FR & ST	16.6	0.0	8.3	2.7	0.0	2.1	2.0	0.0	1.8	7.1	0.0	4.1

**Fig. 4.** Comparison of L2 norms of the classifier weights. Our full method has balanced L2-norms leading to a better balance in classification for old and new classes.

and feature replay (FR). The FRoST without FD and FR results in complete forgetting of the old classes. This is not surprising because without FD the feature extractor has moved far away from the original configuration. Moreover, as the joint classifier weights corresponding to the new classes are only optimized during the NCD stage (due to the disabled loss Eq. (10)), it leads to what is called the task-recency bias, resulting in higher norms for the new classes weights (see Fig. 4). In other words, for any old test sample, the classifier is highly biased to predict the new classes, leading to complete misclassification of old classes. Similar effect is observed when FD is enabled but FR is disabled.

Effectiveness of Self-Training. In the bottom half of Tab. 1 we show the impact of the absence of self-training (ST) on the performance. The FRoST w/o ST causes no interference in the optimization from the new classes and the joint classifier is able to preserve the performance on old classes. This highlights the truly complex nature of balancing the performance of both old and new classes in the class-iNCD setting. This phenomenon is visualized in Fig. 4 through the norms of the weights of the joint classifier where exists large discrepancies between the old versus new classes. Similar conclusions can also be drawn by observing the confusion matrix reported in Fig. 5. Furthermore, when we drop FD and FR along with ST, we notice a further degradation of the performance of old classes, demonstrating the positive impact of FR and FD in not forgetting.

Comparison of our Feature Replay and Distillation with LwF. Here we empirically demonstrate the need of decoupling the LwF objective into FR and FD when it comes to learning a joint classifier for the class-iNCD. As a control experiment, we use the original formulation of LwF (as in Eq. (2)) as a drop-in replacement for our FR and FD. We optimize with the original LwF loss,




Fig. 5. Comparisons of confusion matrix of different methods. Note that, the label IDs of novel classes are re-assigned by our evaluation protocol.

applied both at the softmax and pre-softmax outputs, and found that compared to FRoST the performance on the new classes improves but at the cost of a large drop in the old classes performance. We conjecture that since the weights for the new classes are randomly initialized in $h^{[A]}$ at the start of the NCD stage, the joint classifier as a whole learns at a different rate than the feature extractor g , which is already pre-trained on the old classes. As the LwF loss optimizes both the g and $h^{[A]}$, it causes slow-fast learning interference from the new classes. This is evident from better new classes performance *w.r.t* FRoST. As shown in Tab. 2, indeed adding FR to LwF improves the performance on the old classes (*e.g.*, 6.8% vs 49.9% for Tiny-ImageNet). This again proves the effectiveness of FR and the need of decoupled *not-forgetting* objective in the class-iNCD.

Table 2. Ablation study comparing FRoST with LwF (logits-KD).

IL Methods	CIFAR-10			CIFAR-100			Tiny-ImageNet			Average		
	Old	New	All	Old	New	All	Old	New	All	Old	New	All
FRoST (Ours)	77.4	49.5	63.5	62.5	45.8	59.2	54.4	33.9	52.4	64.8	43.1	58.3
LwF (softmax)	13.6	63.2	38.4	7.4	63.5	18.6	2.1	42.8	6.2	7.7	56.5	21.1
LwF (softmax) + FR	21.4	61.1	41.3	33.3	61.2	38.9	35.3	33.1	35.0	30.0	51.8	38.4
LwF (pre-softmax)	19.4	76.3	47.9	13.6	61.4	23.2	6.8	38.7	10.0	13.3	58.8	27.0
LwF (pre-softmax) + FR	24.8	77.5	51.1	49.3	58.3	51.1	49.9	26.8	47.6	41.3	54.2	49.9

Effect of using Joint and Novel Classifiers. Here we elaborate on the choice of using the joint $h^{[A]}$ and novel $h^{[U]}$ classification heads in FRoST to address class-iNCD. We report the Joint baseline in Tab. 3, where we discard the novel classifier head and solely use the joint classifier for learning the new classes and the old classes. We find that this approach causes hindrance in learning the new classes for CIFAR10 and Tiny-Imagenet data sets because two parts of the same classifier are subject to gradients of different magnitudes, highlighting the need to decouple the learning of two tasks. In the next ablation, we also disable the ST with pseudo-labels that are generated by joint itself and we find that it destabilizes the performance on the new classes. Finally, we construct an ablation where we do not extend $h^{[L]}$ to $h^{[A]}$, but instead use $h^{[L]}$ in conjunction with $h^{[U]}$ and is denoted with Novel w/o ST. We observe that this behaves similarly with the previous ablation analysis of FRoST w/o ST in Tab. 1. Thus, we conclude that having joint and novel heads trained with ST is crucial in class-iNCD.

Table 3. Ablation study on having a single and separated heads for old and new classes. Joint: class-agnostic head; Novel: new classes classifier head.

Classifier Head	CIFAR-10			CIFAR-100			Tiny-ImageNet			Average		
	Old	New	All	Old	New	All	Old	New	All	Old	New	All
Joint + Novel (Ours)	77.4	49.5	63.5	62.5	45.8	59.2	54.4	33.9	52.4	64.8	43.1	58.3
Joint	81.3	41.5	61.4	64.5	46.3	60.9	56.8	8.4	52.0	67.5	32.1	58.1
Joint w/o ST	91.7	0.0	45.8	68.6	29.4	60.7	57.5	0.1	51.7	72.6	9.9	52.8
Novel w/o ST	92.0	0.0	46.0	67.9	32.1	60.7	57.9	0.0	52.1	72.6	10.7	52.9

Table 4. Comparison with state-of-the-art methods in class-iNCD.

Methods	CIFAR-10			CIFAR-100			Tiny-ImageNet			Average		
	Old	New	All	Old	New	All	Old	New	All	Old	New	All
AutoNovel[15]	27.5	3.5	15.5	2.6	15.2	5.1	2.0	26.4	4.5	10.7	15.0	8.4
ResTune[29]	91.7	0.0	45.9	73.8	0.0	59.0	44.3	0.0	39.9	69.9	0.0	48.3
NCL[37]	92.0	1.1	46.5	73.6	10.1	60.9	0.8	6.5	1.4	55.5	5.9	36.3
DTC[14]	64.0	0.0	32.0	55.9	0.0	44.7	35.5	0.0	32.0	51.8	0.0	36.2
FRoST	77.5	49.5	63.4	64.6	45.8	59.2	54.5	33.7	52.3	65.5	39.8	54.9

4.3 Comparison with State-of-the-art Methods

We compare our FRoST with the state-of-the-art NCD methods under the newly proposed class-iNCD setting. We also compare with ResTune [29] which is a recently proposed method for iNCD. As none of these existing methods have been evaluated in the class-iNCD setting, we re-run the baselines and simply modify the evaluation protocol which is described in Sec. 3. We report the results of the NCD [15,37,16], iNCD [29] baselines and FRoST in the Tab. 4. As can be observed, under the class-iNCD all the NCD [15,37,16] fail to obtain a good balance on the old and new classes. Interestingly, while none of these NCD methods use any explicit objectives to prevent forgetting, they tend to predict well the old classes (see column **Old** in Tab. 4) and poor performance on new classes (see column **New** in Tab. 4). When visualizing the confusion matrix in Fig. 5 we found that most of the test samples get classified as old classes due to the old classes classifier having higher norms. As a consequence, this gives the impression that the baselines methods are able to retain performance on old classes. Second, for the above methods, although the new classes performance obtained with the joint head appears to be low, the actual performance of their novel head in the task-aware evaluation is indeed high. We report the breakdown of the novel classes performance in Tab. 5 where, for instance, the column **New-1-N** denotes the task-aware clustering performance of the novel head on the new classes. As can be observed, the new classes classifier of the NCD baselines can indeed learn on the new classes (*e.g.*, 34.2% in NCL vs 32.4% in FRoST).

ResTune, although designed specifically for the iNCD setting, exhibits similar counter-intuitive behaviour with the performance on the old classes dominating the new classes. To investigate this pathology, we inspect into the confusion matrix in Fig. 5 (e) and find that all the samples get predicted to the first five old classes for CIFAR10. In other words, the overall performance reported in ResTune [29] is actually dominated by the old classes performance. We report

Table 5. Comparison with the state-of-the-art methods in the two-step class-iNCD setting where new classes arrive in two episodes, instead of one. New-1-J: new classes performance from joint head at first step, New-1-N: new classes performance from novel head at first step, etc.

Methods	Tiny-ImageNet									
	First Step (180-10)				Second Step (180-10-10)					
	Old	New-1-J	New-1-N	All	Old	New-1-J	New-2-J	New-1-N	New-2-N	All
ResTune[29]	39.7	0.0	38.0	37.6	34.9	0.0	0.0	25.4	42.8	31.4
DTC[14]	38.9	0.0	43.8	36.9	33.4	0.0	0.0	28.0	59.4	30.1
NCL[37]	5.6	0.0	34.2	5.3	1.4	0.0	2.6	21.6	41.6	1.4
FROST	55.2	27.6	32.0	53.8	42.5	34.8	31.2	31.2	46.8	41.6

confusion matrices on bigger data sets in the supplementary material. This shows that the existing evaluation method for iNCD is flawed and our proposed class-iNCD is indeed more meaningful that properly evaluates the effectiveness of a learning algorithm. Contrarily our proposed FROST consistently achieves a good balance in performance in all the tested data sets. This also demonstrates the validity of the components in our proposed FROST. We present a detailed comparison analysis between ResTune and FROST in the supplementary material.

Two-Step Class-iNCD. As done in the class-IL literature [30], we also run experiments on a sequence of novel tasks, which we call as two-step class-iNCD, where 20 novel classes in Tiny-ImageNet are added in two steps, each step dealing with 10 novel classes. We compare our FROST with the baseline methods in Tab. 5 where we show not only the joint classifier head performance (*e.g.*, **New-1-J**), but also from the novel classifier head (*e.g.*, **New-1-N** and **New-2-N**) at each step. As can be seen, for the baseline methods the novel classifier heads can satisfactorily discover the new classes at each step, but when evaluated with the joint head biases the predictions to the old classes. Unlike the baselines, FROST does not suffer from this issue and leads to more balanced predictions.

5 Conclusion

In this work we addressed the novel problem of class-incremental NCD. This task differs from the traditional NCD since we are not only interested in discovering novel classes but also aim to prevent forgetting on the old classes. To prevent this forgetting phenomenon we proposed feature replay and feature-level distillation that is well suited for the class-iNCD. Moreover, to make inference task-agnostic, we propose to maintain a joint classifier that can classify any of the previously seen classes. We train this joint classifier using the pseudo-labels generated by the novel classifier head that is trained with a clustering loss. We compared our method to many relevant works and obtained superior performance on various benchmarks. Given the practical nature of class-iNCD and encouraging results with our FROST, we believe this work will stimulate further research.

Acknowledgements. We thank the funding agencies: EU H2020 projects SPRING (No. 871245) and AI4Media (No. 951911); and the EUREGIO project OLIVER.

References

1. Anderson, B.L.: Can computational goals inform theories of vision? *Topics in Cognitive Science* (2015) [1](#)
2. Buzzega, P., Boschini, M., Porrello, A., Abati, D., Calderara, S.: Dark experience for general continual learning: a strong, simple baseline. In: *NeurIPS (2020)* [3, 4](#)
3. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: *NeurIPS (2020)* [1](#)
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *ICCV (2021)* [1](#)
5. Castro, F.M., Marín-Jiménez, M.J., Guil, N., Schmid, C., Alahari, K.: End-to-end incremental learning. In: *ECCV (2018)* [4](#)
6. Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: *ECCV (2018)* [4](#)
7. Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P.K., Torr, P.H., Ranzato, M.: Continual learning with tiny episodic memories. In: *ICML (2019)* [3, 4](#)
8. Delange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T.: A continual learning survey: Defying forgetting in classification tasks. *TPAMI (2021)* [2, 8](#)
9. Dhar, P., Singh, R.V., Peng, K.C., Wu, Z., Chellappa, R.: Learning without memorizing. In: *CVPR (2019)* [4](#)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR (2021)* [1](#)
11. Fini, E., Sangineto, E., Lathuilière, S., Zhong, Z., Nabi, M., Ricci, E.: A unified objective for novel class discovery. In: *ICCV (2021)* [1, 2, 3, 5](#)
12. French, R.M.: Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences* (1999) [1](#)
13. Goodfellow, I.J., Mirza, M., Xiao, D., Courville, A., Bengio, Y.: An empirical investigation of catastrophic forgetting in gradient-based neural networks. In: *ICLR (2014)* [4](#)
14. Han, K., Vedaldi, A., Zisserman, A.: Learning to discover novel visual categories via deep transfer clustering. In: *ICCV (2019)* [13, 14](#)
15. Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Automatically discovering and learning new visual categories with ranking statistics. In: *ICLR (2020)* [1, 2, 3, 6, 8, 10, 13](#)
16. Han, K., Vedaldi, A., Zisserman, A.: Learning to discover novel visual categories via deep transfer clustering. In: *ICCV (2019)* [1, 2, 3, 4, 13](#)
17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR (2016)* [1, 10](#)
18. Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. In: *NeurIPS Workshop (2014)* [6, 7](#)
19. Hsu, Y.C., Lv, Z., Kira, Z.: Learning to cluster in order to transfer across domains and tasks. In: *ICLR (2018)* [2, 3](#)
20. Hsu, Y.C., Lv, Z., Schlosser, J., Odom, P., Kira, Z.: Multi-class classification without multi-class labels. In: *ICLR (2019)* [2, 3](#)

21. Jia, X., Han, K., Zhu, Y., Green, B.: Joint representation learning and novel category discovery on single-and multi-modal data. In: ICCV (2021) [3](#)
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. In: NAS (2017) [4](#)
23. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. University of Tronto (2009) [10](#)
24. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NeurIPS (2012) [1](#)
25. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly (1955) [5](#)
26. Le, Y., Yang, X.: Tiny imagenet visual recognition challenge. CS 231N (2015) [10](#)
27. Li, Z., Hoiem, D.: Learning without forgetting. TPAMI (2017) [2, 4, 6, 7](#)
28. Liang, J., Hu, D., Feng, J.: Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In: ICML (2020) [2](#)
29. Liu, Y., Tuytelaars, T.: Residual tuning: Toward novel category discovery without labels. TNNLS (2022) [2, 3, 4, 5, 7, 10, 13, 14](#)
30. 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. arXiv preprint arXiv:2010.15277 (2020) [4, 14](#)
31. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: CVPR (2017) [3, 4](#)
32. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Fu, Y.: Large scale incremental learning. In: CVPR (2019) [4](#)
33. Xie, J., Girshick, R., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: ICML (2016) [3](#)
34. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: ICML (2017) [4](#)
35. Zhao, Y., Zhong, Z., Sebe, N., Lee, G.H.: Novel class discovery in semantic segmentation. In: CVPR (2022) [3](#)
36. Zheng, Z., Yang, Y.: Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. IJCV (2021) [2](#)
37. Zhong, Z., Fini, E., Roy, S., Luo, Z., Ricci, E., Sebe, N.: Neighborhood contrastive learning for novel class discovery. In: CVPR (2021) [1, 2, 3, 5, 10, 13, 14](#)
38. 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. In: CVPR (2021) [1, 2, 3](#)