arXiv:2304.14310v1 [cs.CV] 27 Apr 2023

Incremental Generalized Category Discovery

Bingchen Zhao Oisin Mac Aodha
University of Edinburgh
https://bzhao.me/iNatIGCD

Abstract

We explore the problem of Incremental Generalized Category Discovery (IGCD). This is a challenging category-incremental learning setting where the goal is to develop models that can correctly categorize images from previously seen categories, in addition to discovering novel ones. Learning is performed over a series of time steps where the model obtains new labeled and unlabeled data, and discards old data, at each iteration. The difficulty of the problem is compounded in our generalized setting as the unlabeled data can contain images from categories that may or may not have been observed before. We present a new method for IGCD which combines non-parametric categorization with efficient image sampling to mitigate catastrophic forgetting. To quantify performance, we propose a new benchmark dataset named iNatIGCD that is motivated by a real-world fine-grained visual categorization task. In our experiments we outperform existing related methods.

1. Introduction

The wealth and complexity of visual information potentially observable by artificial systems deployed in the real world vastly exceeds the comparative simplicity of our carefully curated benchmark vision datasets. To operate safely and reliably in challenging environments, these systems need to be able to correctly recognize previously learned concepts, not confuse these known concepts with novel ones, and be able to differentiate novel concepts so that they can be grouped and efficiently learned. As humans, we excel at this type of flexible learning in such dynamic settings [1, 29], and it is clear that we need to endow our artificial systems with such desirable properties.

In the context of visual categorization, there is a rich body of work that has moved beyond the traditional supervised setting into more open-ended learning paradigms. For example, in semi-supervised learning, in addition to labeled data, during training we also have access to unlabeled data depicting the same categories [59]. Other work has attempted to address more complex settings such as determining if an image contains a previously observed, or instead novel category (i.e. open-set recognition) [45], and learning to group images from novel categories (i.e. novel category discovery) [16, 49]. Recently, a more realistic and challenging setting termed Generalized Category Discovery (GCD) has been proposed [6, 51]. Here, unlabeled images can be from either previously seen *or* from novel categories and the task is to develop a model that can both classify the previously seen categories and also discover novel ones.

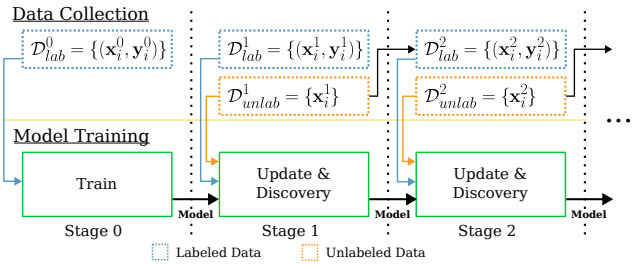

Figure 1: Overview of our incremental generalized category discovery setting. At each time step the model has access to labeled and unlabeled data and must simultaneously classify existing categories and discover new ones. As time progress, all previous data is no longer directly available.

In this work, we go one step further by exploring the generalized setting along the temporal, i.e. category-incremental, dimension [48, 39]. We refer to this as **I**ncremental **G**eneralized **C**ategory **D**iscovery (IGCD). Here, learning progress over a series of stages (representing time steps), where at each stage we acquire a new set of unlabeled data that contains images from categories that may, or may not, have been previously observed (see Fig. 1). At each stage, the goal when updating the model is to maintain performance on the previously observed categories, while also discovering novel ones. This must be achieved under the constraint that at each new stage the labeled data from earlier stages is no longer available for training.

A small number of recent works have started to investigate incremental learning in the context of category discovery. For example, when only novel categories are present at each subsequent stage in the case of [42, 23, 32] or in the generalized setting in [56] where the unlabeled data at each

each stage contains both old and novel categories. However, in contrast to [56], in our IGCD setting we do not make the simplifying assumption that the unlabeled data at each stage contains *all* the categories from the previous stages. More realistically, unlabeled data is not guaranteed to contain all previously seen categories which makes the task even harder for the incremental learner. We show that combining a state-of-the-art GCD method with incremental learning techniques can work, but results in forgetting as time progresses. We address these issues via a novel IGCD method that combines a non-parametric classifier with a density-based sampling mechanism that efficiently enables the selection of informative examples for both classification and past memorization.

Benchmarking in the incremental discovery setting to date has largely been restricted to artificial category-based splits of conventional image categorization datasets [27, 53, 26, 35]. To encourage future progress on IGCD, we also present a new dataset called iNatIGCD. iNatIGCD more faithfully simulates a real-world fine-grained incremental learning setting. It is motivated by a real-world visual categorization use-case whereby the categories present at different stages are selected using spatio-temporal metadata from the community science platform iNaturalist [22]. As a result, iNatIGCD naturally includes challenges such as appearance shifts, in addition to both fine and coarse-grained differences between old and novel categories present in images sourced from different geographical regions.

In summary, we present the following contributions: (i) A new approach for IGCD that combines non-parametric classification with efficient incremental learning via density-based support set selection. (ii) iNatIGCD, a new in-the-wild dataset for benchmarking IGCD that features a multi-stage training split motivated by a real-world fine-grained visual categorization use-case. (iii) A thorough evaluation on IGCD where we outperform recent methods.

2. Related Work

Here we review existing work in semi-supervised learning and both standard and incremental category discovery.

Semi-Supervised Learning. Conventional semi-supervised learning assumes that we have access to both labeled and unlabeled data at training time [59, 8, 38]. Many works have been proposed to tackle this task using pseudo-labeling [40], consistency regularization [4, 47, 46, 28], density-based label propagation [30], or non-parametric categorization [3]. It is typically assumed that the unlabeled data contains instances from the *same* categories that are present in the labeled set. Recent works have extended semi-supervised learning by removing the assumption that the categories in the unlabeled and labeled sets are the same [44, 21, 55], but their focus is still on the performance on the labeled categories and they do not evaluate the clustering accuracy on the novel categories in the unlabeled set.

Category Discovery. In contrast to semi-supervised learning, Novel Category Discovery (NCD) [16] addresses an alternative setting whereby there is no overlap between the categories in the labeled and unlabeled sets. Here, the goal is to automatically discover the novel categories in the unlabeled data. This can be viewed as a semi-supervised clustering problem [19, 20, 16].

Recently, Generalized Category Discovery (GCD), a more realistic and challenging version of the discovery problem has been proposed [6, 51]. In this setting, the unlabeled data can consist of images from *both* seen and unseen (i.e. novel) categories. Earlier works showed that self-supervised pretraining can aid category discovery [15]. Starting from a backbone initialized using self-supervised pretraining [9], ORCA [6] assumes the number of novel categories is known and proposed a three-component loss to train a deep classifier. The loss is comprised of a supervised component, a pairwise loss that uses high confident pseudo-labels to enforce that similar unlabeled instances are grouped together, and a regularization term to ensure that all instances are not assigned to the same category. GCD [51] also make use of self-supervised pretraining using [7], but foregoes the need for pseudo-labels by instead using a clustering-based approach. They perform a contrastive finetuning step using image pairs from the same category for the labeled data and augmented pairs for the unlabeled data. Finally, they perform unsupervised clustering using k-means to assign unlabeled images to categories.

Building on GCD [51], SimGCD [54] investigated the impact of different design choices on downstream performance. Their final approach does not use an explicit clustering step but instead makes use of learned category prototypes inspired by self-distillation methods [2], resulting in improved performance compared to GCD. MIB [11] employs a similar pretraining phase to GCD to train their feature extractor, in addition to a cross-entropy and conditional entropy loss for the labeled and unlabeled data respectively. We take inspiration from these methods, but explore a different setting of the discovery problem, that of IGCD.

Incremental Category Discovery. Category-incremental learning is a learning setting whereby the number of categories in a dataset increases over time, e.g. over a set of discrete learning stages. The challenge in the incremental setting is that it is not possible to store all of the previously observed training data from earlier stages during each subsequent learning stage. Some of the difficulties that arise from this setting include catastrophic forgetting [36, 25] (i.e. a dramatic decrease in performance on old categories as a result of training on new ones) and susceptibility to distributional shifts in the data. Multiple different approaches have been explored in the literature to address this, includ-

ing storing past training examples in a replay buffer [41] or storing distilled exemplars [39, 5]. Taxonomic variants of the problem have also been explored [31, 10]. For a recent survey on the topic we refer readers to [58].

Most relevant to our work are a recent set of approaches that explore category discovery in the incremental/continuous setting, e.g. [42, 23, 56, 32]. [42] proposed a one-stage category-incremental setting termed class-iNCD. Here labeled data is initially available to learn a representation, then discarded, and finally a set of unlabeled data containing only novel categories is provided. Their goal is to train a model that performs well on both sets of categories. However, here there is no overlap between the categories in the labeled and unlabeled sets. This setting has also been explored in [33], where their focus was only on the category discovery performance. FRoST [42] retains feature prototypes learned from labeled data which are replayed during the discovery phase to prevent forgetting of the old categories. NCDwF [23] also explore the same setting and propose a method that uses pseudo-latent supervision, feature distillation, and a mutual information-based regularizer to maintain performance on the discarded labeled categories and assist in discovering novel ones. In the MSc-iNCD setting explored in [32], the model only obtains unlabeled data at each time step which is assumed to come from previously unseen (i.e. novel) categories.

Grow and Merge (GM) [56] was recently introduced and applied to several different incremental discovery settings. It performs two key phases at each learning stage. First in the growing phases, it performs novelty detection to separate novel from seen categories and then trains a dynamic network to perform NCD. Then in the merging phases, it combines the newly discovered categories with the previously know ones into a single model. Later, we compare to GM and show superior performance.

With the exception of GM, who only perform one related experiment on CIFAR-100 [27], other existing methods assume that unlabeled data can *only* contain novel categories. In contrast, we explore the more challenging incremental *generalized* setting where, at each stage, new data can come from *either* previously seen or novel categories. Unlike [56], in our setting, all previously seen categories are not guaranteed to be present in each subsequent stage.

Benchmarking Category Discovery. There are several common datasets used to evaluate category discovery methods. Earlier works [16, 15] typically use standard image categorization datasets like CIFAR10/100 [27] and ImageNet [43] by creating artificial splits for discovery evaluation. [57] argue that more challenging fine-grained datasets like CUB [53] are more suitable for evaluating discovery performance as the labeled and unlabeled categories share more visual similarity. Recently, the Semantic-Shift Benchmark (SSB) [52, 51] was proposed to better evaluate the task of detecting semantic novelty. SSB uses the existing CUB [53], Stanford Cars [26], and FGVC-Aircraft [35] datasets where the category splits are designed to have a clear 'axes of semantic variation' as well as a coherent definition of categories in the labeled and unlabeled sets.

Recent work in the incremental discovery setting [42, 23, 56] evaluate performance using artificial data splits on CIFAR-10/100 [42, 56, 23], TinyImageNet [42], ImageNet [23], or the fine-grained SSB datasets [56]. However, these artificial data splits may not reflect the real world performance of IGCD methods. For example, it has been shown that the category splits in the labeled dataset play an important role NCD performance [13]. To address this gap, we introduce a new 'real-world' data split derived from the publicly available fine-grained iNaturalist species categorization dataset [50]. We make it applicable to the incremental setting by leveraging metadata, such as the date and location of each observation. From this, we are able to better simulate real-world IGCD by creating data and category splits that mimic those that naturally arise over multiple time steps in the real-world.

3. The iNatIGCD Benchmark

Here we outline iNatIGCD, our new benchmark for IGCD. Our dataset is based on the large-scale fine-grained iNat2021 visual categorization dataset [50], which contains images from 10,000 different categories of plant and animal species sourced from the community science platform iNaturalist [22]. Each of the images posted to iNaturalist contains the precise capture date and time, along with the geographical location as metadata [22]. This creates a natural stream of data over time that contains category distribution shifts between observations from around the globe as the iNaturalist community grows across different regions.

iNatIGCD is motivated by the scenario where one can develop an initial image classifier from the currently available labeled data, which only contains images and labels captured before a certain time and is restricted to a certain region (e.g. a specific continent). As time progresses, new unlabeled data is obtained which can also include images other regions (e.g. a different continent). This scenario is close to the real-world situation of the iNaturalist platform where we have access to labeled image observations at a given time point, but we cannot easily obtain reliable labels for new observations that are shared each day as labeling images takes time. We can view the existing collection of observations as the labeled dataset and the new observations, which can contain a combination of old and new categories, as the unlabeled dataset. Similar to the common assumption in continual learning [39, 5], all the previous labeled data can be difficult, or even impossible, to store during learning. Thus as time progresses, we assume that we do *not* have access to all the past labeled data when retraining/updating the model.

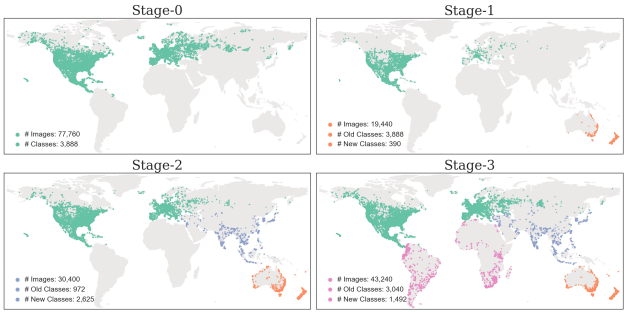

Figure 2: A visualization of the location of the data in each stage of our iNatIGCD dataset. We also report the number of old and new categories for each of the stages.

Our new IGCD benchmark is structured as follows: At each time step, we have a labeled dataset and an unlabeled one. Our goal is to correctly discover any novel categories in the unlabeled dataset while also classifying examples belonging to known categories. When we progress to a new time step, the labeled dataset from the previous step becomes unavailable. However, we now have a new labeled dataset, which is actually the previous unlabeled dataset, along with a new unlabeled dataset which is used for category discovery. This simulates the real-world process as it takes time for the iNaturalist community to reach a consensus regarding which category is present in a newly uploaded image. See Fig. 1 for an illustration of the process.

Dataset Construction. To generate our data splits, we leverage metadata from the iNat2021 dataset [50]. We first sort and group all the images according to the time when they were captured. To make the data more evenly distributed, we cluster them into four temporal stages: 2008-2016, 2017, 2018, and 2019. We further split the data according to the location of the observation, only including data from North America and Europe in the first stage and then including Oceania, Asia, and finally Africa and South America in the second, third, and fourth stages respectively. This simulates the growth of the species recognition community on iNaturalist from continent to continent over time.

Our new iNatIGCD dataset is summarized in Fig. 2. It poses several unique challenges for category discovery methods: (i) we have an order of magnitude more finer-grained categories in each of the incremental stages to challenge the ability of category discovery methods, (ii) at each stage, we have more images that require categorization, which can give a more reliable measure of the discovery performance of a model, and (iii) the categories that appear in each stage are based on the natural distribution of species rather than an artificial split as in previous benchmarks.

4. Method

In this section we present a new approach targeting the IGCD setting. At each stage t (i.e. time step), we have two datasets to train the model on, $\mathcal{D}^t_{lab}$ and $\mathcal{D}^t_{unlab}$. Our goal is to learn a classifier that can correctly classify and discover categories in $\mathcal{D}^t_{lab}$ and $\mathcal{D}^t_{unlab}$, in addition to maintaining its performance on the categories in the current stage when we advance to later stages.

In the IGCD setting, it is assumed that it is not possible to store all the previously observed data. As a result, the model must learn incrementally. A common incremental learning technique is to avail of the concept of example replay [39, 5]. This involves saving a few examples of each category and then 'replaying' them to the model during training in later stages. However, there are challenges associated with applying replay approaches in our IGCD setting due to two primary issues: (i) we have novel categories in the unlabeled dataset in IGCD, thus it is not straightforward to select examples for these novel categories as previous works typically assume full supervision, and (ii) SoTA GCD methods use parametric classifiers [54, 14], which we show in experiments can overfit to the few replay examples. We tackle these technical challenges by using a non-parametric classifier combined with a density-based example selection mechanism which we introduce below.

4.1. Problem Setting

Formally, we denote the labeled dataset at the initial stage as $\mathcal{D}^0_{lab} = \{(\boldsymbol{x}^0_i, \boldsymbol{y}^0_i)\}$. This is used to train an initial model to recognize the set of categories $\mathcal{C}^0 = \{1, 2, \ldots, K^0\}$. For incremental learning at a later stage t, we have two datasets, a labeled dataset $\mathcal{D}^t_{lab} = \{(\boldsymbol{x}^t_i, \boldsymbol{y}^t_i)\}$ and an unlabeled dataset $\mathcal{D}^t_{unlab} = \{\boldsymbol{x}^t_i\}$. We denote the category sets in $\mathcal{D}^t_{lab}$ and $\mathcal{D}^t_{unlab}$ as $\mathcal{C}^t_{lab}$ and $\mathcal{C}^t_{unlab}$. In our generalized setting, $\mathcal{C}^t_{lab}$ and $\mathcal{C}^t_{unlab}$ can be different and may only have partial overlap. When we progresses to stage $t+1$, $\mathcal{D}^t_{unlab}$ will receive its labels and then be used as $\mathcal{D}^{t+1}_{lab}$ for the next stage. The goal is to learn a model that can not only discover novel categories in the unlabeled dataset $\mathcal{D}^t_{unlab}$ at stage t, but also maintain the ability to recognize the categories in $\mathcal{C}^0_{lab}, \mathcal{C}^1_{lab}, \ldots, \mathcal{C}^t_{lab}$. Similar to the category-incremental learning setting [39, 5], one main added challenge of our setting compared to static GCD is that the model can exhibit catastrophic forgetting [41] on the categories that are not present in $\mathcal{C}^t_{lab}$ and $\mathcal{C}^t_{unlab}$.

Note, our setting is related, but different from the Mixed Incremental (MI) one in [56]. In contrast to us, [56] assume that only the unlabeled data $\mathcal{D}^t_{unlab}$ at stage t is available for training the model. In our experiments, we also compare models that are only trained with $\mathcal{D}^t_{unlab}$. Furthermore, in the MI setting, the category set $\mathcal{C}^t_{unlab}$ always includes *all* previously seen categories. In our more difficult and realistic setting, the category set at each time step may only have partial overlap with the previous category sets.

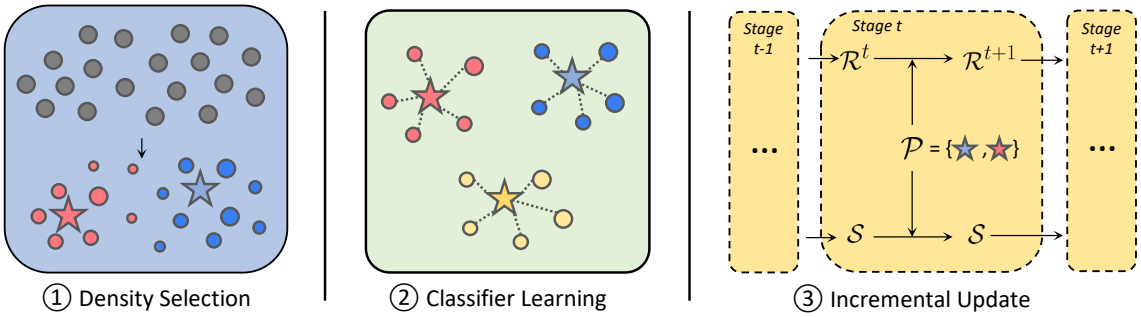

Figure 3: Overview of our method. (1) Illustration of the density calculation process. We denote the density of each data point by its size in the lower part. The density peaks, i.e. the data point whose density is higher than its neighbors, are illustrated with a star. (2) The learning of classifiers, we update representations and classifiers using the loss defined in Eq. (7). (3) The incremental update procedure for stage t. When proceeding to stage $t+1$, the density peaks $\mathcal{P}$ are added to $\mathcal{R}^{t+1}$ and $\mathcal{S}$.

4.2. Classifier Learning

For each input image $\boldsymbol{x}_i$, we use a neural network f to extract its corresponding feature representation $\boldsymbol{h}_i = f(\boldsymbol{x}_i)$. From this, we use a non-parametric classifier g to assign the image to a category probability vector $\boldsymbol{p}_i = g(\boldsymbol{h}_i)$, where $\sum_c \boldsymbol{p}_{i,c} = 1$. For f to learn an informative representation, we adopt a supervised contrastive loss $\mathcal{L}_{\text{SupCon}}$ form [24] for learning from $\mathcal{D}_{lab}$ and use a self-supervised contrastive loss $\mathcal{L}_{\text{SelfCon}}$ from [9] for learning from $\mathcal{D}_{unlab}$. These two losses are combined via a balancing factor λ_{rep},

$$\mathcal{L}_{\text{rep}} = \lambda_{\text{rep}}\mathcal{L}_{\text{SupCon}} + (1-\lambda_{\text{rep}})\mathcal{L}_{\text{SelfCon}}. \tag{1}$$

We use a non-parametric Soft Nearest-neighbor (SNN) classifier for g that computes the distribution over category labels for an input image $\boldsymbol{x}_i$,

$$\boldsymbol{p}_i = \texttt{SNN}(\boldsymbol{h}_i, \mathcal{S}, \tau) = \sum_{k=1}^{N^{\mathcal{S}}} \frac{\exp(\boldsymbol{h}_i \cdot \boldsymbol{h}_k^l/\tau)}{\sum_{\boldsymbol{h}_b^l \in \mathcal{S}} \exp(\boldsymbol{h}_i \cdot \boldsymbol{h}_b^l/\tau)} \boldsymbol{y}_k. \tag{2}$$

$\mathcal{S} = \{(\boldsymbol{h}_k^l, \boldsymbol{y}_k)\}$ is a support set containing features of all support samples $\boldsymbol{h}_k^l$ and their corresponding one-hot labels $\boldsymbol{y}_k$, $N^{\mathcal{S}} = |\mathcal{S}|$ is the number of support samples in $\mathcal{S}$, and $K^{\mathcal{S}}$ is the number of categories in $\mathcal{S}$. $\boldsymbol{h}_i$ is the representation of $\boldsymbol{x}_i$ from f, and τ is a temperature parameter controlling the sharpness of the prediction $\boldsymbol{p}_i$. The support samples are representative examples of each of the categories and are used to softly assign a label to the input example $\boldsymbol{h}_i$ by averaging all the one-hot labels $\boldsymbol{y}_k$ in $\mathcal{S}$ weighted by the softmax similarity. For the labeled categories, we randomly select the support examples. Later we discuss the selection of support examples for the novel categories.

Importantly, as our classifier is non-parametric, it is easy to extend it so that it can classify more categories by simply adding additional examples from the new categories to the support set. At test time, we just use the saved support examples to perform categorization for an input image by assigning the labels via measuring the distance between the input image and support examples. The classifier is trained with two cross-entropy losses $\mathcal{L}_{\text{ce}}$, and makes use of an entropy-regularizer,

$$\mathcal{L}_l = \frac{1}{|\mathcal{B}^l|} \sum_{i \in \mathcal{B}^l} \mathcal{L}_{\text{ce}}(\boldsymbol{y}_i^l, \hat{\boldsymbol{p}}_i^l), \tag{3}$$

$$\mathcal{L}_u = \frac{1}{|\mathcal{B}^u|} \sum_{i \in \mathcal{B}^u} \mathcal{L}_{\text{ce}}(\tilde{\boldsymbol{p}}_i^u, \hat{\boldsymbol{p}}_i^u) - \epsilon H(\overline{\boldsymbol{p}}), \tag{4}$$

where $\mathcal{B}^l$ and $\mathcal{B}^u$ are batches of indices from the labeled and unlabeled datasets, $\boldsymbol{y}_i^l$ is the ground-truth label for the example $\boldsymbol{x}_i^l$, and $\tilde{\boldsymbol{p}}_i^u$ and $\hat{\boldsymbol{p}}_i^u$ are the predictions for two augmented views of the same input example $\boldsymbol{x}_i^u$ with different temperatures. These two losses together form the loss function we use for classifier learning,

$$\mathcal{L}_{\text{cls}} = \lambda_{\text{cls}}\mathcal{L}_l + (1-\lambda_{\text{cls}})\mathcal{L}_u, \tag{5}$$

where λ_{cls} is a weighting factor. The entropy regularization term, $H(\overline{\boldsymbol{p}}) = -\sum \overline{\boldsymbol{p}} \log \overline{\boldsymbol{p}}$, regularizes the mean prediction $\overline{\boldsymbol{p}}$ over a mini-batch which is computed as

$$\overline{\boldsymbol{p}} = \frac{1}{2|\mathcal{B}^u|} \sum_{i \in \mathcal{B}^u} (\tilde{\boldsymbol{p}}_i^u + \hat{\boldsymbol{p}}_i^u). \tag{6}$$

This above entropy term has been adopted in previous works in GCD [54] and semi-supervised learning [3] to calibrate predicted category distribution and to avoid empty clusters, as observed in [51, 54]. The overall loss for training the model is

$$\mathcal{L} = \mathcal{L}_{\text{rep}} + \mathcal{L}_{\text{cls}}. \tag{7}$$

4.3. Support Sample Selection

Our SNN classifier function is similar to the one defined in the semi-supervised approach of [3]. However, in its original form, it cannot be used to learn from novel categories from the unlabeled data if they do not already appear in the

labeled set. We extend this approach via a novel modification to make it applicable to the generalized category discovery setting. Specifically, for the potentially novel categories in the unlabeled dataset, we propose a new density-based selection mechanism to select a subset of examples from the unlabeled dataset and then pseudo-label them with g to use them as support samples in S.

We start by estimating the density d_i of a sample $\boldsymbol{x}_i^u$ based on its k-nearest neighbors, such that

$$d_i = \frac{1}{K} \sum_{j \in \mathcal{N}_{\boldsymbol{x}_i^u}}^{K} \frac{\boldsymbol{h}_i^u \cdot \boldsymbol{h}_j^u}{\|\boldsymbol{h}_i^u\|_2 \|\boldsymbol{h}_j^u\|_2}. \tag{8}$$

Here we iterate through the top-K neighbors and average the cosine similarities of the feature representations, where $\mathcal{N}_{\boldsymbol{x}_i^u}$ contains the indices of the K nearest neighbors of $\boldsymbol{x}_i^u$. Intuitively, a larger density value d_i indicates that the corresponding prediction for $\boldsymbol{x}_i^u$ is a more reliable pseudo-label as it indicates that $\boldsymbol{x}_i^u$ is similar to its neighbors in the learned representation space of f, see [30]. Inspired by this, we propose a method for automatically selecting reliable samples from $\mathcal{D}_{unlab}$ using this density definition.

First, we select the density peaks from the dataset which are the points whose density is higher than all other k-nearest neighbors around it [30]. Specifically, the set of density peaks can be defined as

$$\mathcal{P} = \{\boldsymbol{x}_i^u | \forall k \in \mathcal{N}_{\boldsymbol{x}_i^u}^K, d_i > d_k\}, \tag{9}$$

where $\mathcal{P}$ is a set of selected images $\boldsymbol{x}_i$. Second, it is not guaranteed that there is only one density peak per class. To address this, we use an intersection-over-union (IoU) score on the nearest neighbor set between two density peaks to measure their similarity. We then remove redundant peaks using a similar procedure to non-maximum-suppression [37].

Thus, for two density peaks $\boldsymbol{x}_i^u$ and $\boldsymbol{x}_j^u$ in $\mathcal{P}$, we measure their similarity by calculating the IoU score of their K^d nearest neighbor set $\mathcal{N}_{\boldsymbol{x}_i^u}^{K^d}$ and $\mathcal{N}_{\boldsymbol{x}_j^u}^{K^d}$. If the IoU score is higher than a threshold T, we only keep the peak with the higher density value. The selected density peaks are used to pseudo-label their close neighbors based on distance, and then the density peaks and their closer neighbors are added as support samples for novel categories to $\mathcal{S}$. For the labeled categories, we select a number of support examples from the labeled dataset using the same density selection mechanism to serve as their support samples. This procedure is outlined in detail in the supplementary material.

The benefit of our proposed method is that we automatically have an estimate of the number of novel categories in the unlabeled set based on the number of selected density peaks. This is in contrast to previous methods for estimating the number of categories which make use of the k-means algorithm, e.g. [16, 51]. Our approach is more closely integrated with the representation learning step and is also more computationally efficient.

4.4. Incremental Update

Earlier we introduced our method for learning a classifier that leverages non-parametric category assignments. Now we discuss how to extend this model to the *incremental* setting. To do this, we focus on learning between a pair of stages, i.e. from stage t to $t+1$.

After training on $\mathcal{D}_{lab}^t$ and $\mathcal{D}_{unlab}^t$ with the objectives defined in Sec. 4.2, we have a model that can classify the labeled and unlabeled categories, $\mathcal{C}_{lab}^t$ and $\mathcal{C}_{unlab}^t$. To handle subsequent stages, we draw inspiration from the classic category-incremental learning method iCaRL [39]. Specifically, we store a set of representative examples $(\boldsymbol{x}_i, \boldsymbol{y}_i)$ to a memory buffer $\mathcal{R}^t$ for replay to be used for training at later stages. Note, this is not to be confused with the support set S used by our non-parametric classifier defined earlier.

Different from iCaRL which selects examples using category centroids, we use our proposed density selection mechanism described in Sec. 4.3 to select the examples to save to $\mathcal{R}^t$. Thus at stage $t+1$, we concatenate the training set $\mathcal{D}_{lab}^{t+1}$ with the memory buffer $\mathcal{R}^t$ and then optimize the training objectives in Sec. 4.2. The memory buffer created at stage $t+1$ is also concatenated with $\mathcal{R}^t$ to ensure that the knowledge from all previous stages is maintained.

5. Experiments

5.1. Implementation Details

Model Training. We use a ResNet-18 [18] initialized using MoCo [17] on ImageNet-1k as our feature extractor f. We freeze its parameters up to the last residual block to prevent the model from overfitting on the labeled data as in [15, 57]. We train for 100 epochs on the labeled data from the initial stage $\mathcal{D}_{lab}^0$, and then train on each subsequent stage t with $\mathcal{D}_{lab}^t$ and $\mathcal{D}_{unlab}^t$ for 40 epochs each. We provide additional vision transformer [12] results and details of the training settings in the supplementary material.

Evaluation Metrics. We use our iNatIGCD dataset to evaluate the performance of different IGCD methods, using two different settings: (i) IGCD-l, where we assume the labels of $\mathcal{D}_{unlab}^t$ will become available at the end of stage t and thus can be used as $\mathcal{D}_{lab}^{t+1}$, and (ii) IGCD-u, where we assume no data is labeled in any of the incremental stages, and at stage t only $\mathcal{D}_{unlab}^t$ is available to train the model. We remove the supervised losses accordingly in IGCD-u.

For our experiments on iNatIGCD, we report the clustering accuracy for categories in $\mathcal{C}_{lab}^t$ and $\mathcal{C}_{unlab}^t$ at stage t as 'Old' and 'New', and 'All' as the performance on both. We also report the accuracy on the categories in previous stages $0, 1, \ldots, t-1$ that do not appear in the current stage as 'S-t'. As in [56], we report maximum forgetting $\mathcal{M}_f$, which is the

Methods	Stage-0	Stage-1			Stage-2				Stage-3					Overall	
	All	All	Old	New	All	Old	New	S-0	All	Old	New	S-1	S-0	$M_f \downarrow$	M_d
Supervised upper-bound	42.3	62.8	68.7	59.8	65.8	71.5	62.3	48.2	68.5	72.5	64.3	58.4	49.7	-	-
SimGCD [54]	42.3	24.1	**35.8**	10.5	25.6	**36.8**	18.4	2.5	22.4	**36.7**	16.4	5.4	1.3	41.0	24.5
SimGCD + iCaRL [39]	41.5	23.7	33.6	11.9	23.5	34.8	12.7	9.4	**23.5**	35.8	17.5	11.3	8.6	32.9	27.9
FRoST [42]	42.3	20.5	24.6	9.4	12.4	24.1	7.8	7.3	13.5	18.5	7.6	6.5	4.2	38.1	19.7
GM [56]	42.3	18.7	26.7	8.6	19.5	28.7	10.4	14.0	16.5	25.8	13.5	**16.8**	12.3	30.0	20.1
Ours	41.8	**25.6**	34.6	**14.5**	**25.7**	35.0	**21.5**	**16.4**	22.8	35.9	**17.8**	16.4	**14.2**	**27.6**	**28.4**

Table 1: Results on iNatIGCD in the IGCD-l setting (i.e. where labels are available at the end of each stage). Higher numbers are better, with the exception of M_f where lower is better.

maximum difference between the clustering accuracy from the stage 0 categories and any later stage t. We also report final discovery $\mathcal{M}_d$, which is the final clustering accuracy of the model on all the categories from all stages.

5.2. Results

Here we describe the results of relevant methods on different IGCD benchmarks. We include comparisons to GM [56], a recent IGCD method. We also compare to FRoST [42] which is a one-stage method that can only handle one incremental stage. We extend it to the multi-stage setting by repeating the second discovery phase of FRoST [42] for each new stage. We also compare to SimGCD [54], a recent SoTA GCD method and apply it to IGCD by rerunning it on the data from each new stage. To enable it to maintain performance on previous categories, we combine it with a memory buffer as in iCaRL [39]. Finally, to establish a performance upper bound on iNatIGCD, we train a fully supervised model on all the data and labels at each stage t (including all previous stages $t-1, \dots, 0$). Additional results are in the supplementary material.

5.2.1 iNatIGCD Results

IGCD-l Setting. We first present results for the incremental setting with labels in Tab. 1. Although the non-incremental SimGCD [54] achieves comparable performance at recognizing labeled categories ('Old') and discovering novel categories ('New') at each stage, it fails to maintain good performance on the categories it has seen before but that are not present in the current stage (see 'S-0' and 'S-1'). Combining SimGCD with iCaRL [39] helps alleviate this forgetting problem, which we see as an increase in 'S-0' and 'S-1' scores compared to SimGCD alone. However, the performance on 'Old' categories, which have labeled data during training, drops for SimGCD + iCaRL. We argue that this is because SimGCD employs a parametric classifier, but iCaRL [39] performs non-parametric categorization at test time. We can see from the Stage-0 performance of SimGCD + iCaRL that this mismatch between training and testing results in a performance drop even in the initial fully supervised scenario. In contrast, our SNN classifier achieves a balance between preventing the category information from previous stages from being forgotten ('S-0' and 'S-1') as well as obtaining good performance on the categories it is currently trained on ('Old' and 'New').

Compared to GM [56], which is designed to tackle discovery and forgetting simultaneously, our proposed method shows a clear advantage, with improved discovery performance (see $\mathcal{M}_d$) and reduced forgetting (see $\mathcal{M}_f$). We also outperform the recent FRoST [42] method. Notably, we always perform best on the novel categories at each stage (see 'New'). We speculate that this is likely because GM [56] and FRoST [42] employ a pair-wise objective for learning to cluster novel categories. However in iNatIGCD, the number of categories is high, thus there will be far more negative pairs than positive ones to train their methods, resulting in the degraded performance.

IGCD-u Setting. In Tab. 2 we present results for the incremental setting where $\mathcal{D}^t_{lab}$ is not available at each stage t. Thus we report performance on 'New' categories in $\mathcal{D}^t_{unlab}$ at each stage t to compare a model's ability to discover new categories. We also report the performance on S-0, S-1, and S-2 to evaluate the forgetting of previously learned categories. Due to the lack of supervision from $\mathcal{D}^t_{lab}$, the performance of all models is reduced, yet our proposed approach achieves superior results in almost all cases.

5.2.2 Mixed Incremental Results

In Tab. 3 we present results on the Mixed Incremental (MI) scenario proposed in [56]. At each incremental stage t, the model is trained on a dataset $\mathcal{D}^t_{unlab}$ of unlabeled examples which contains *both* novel categories and the categories the model has already learned before. The goal is to classify both novel and seen categories at the same time. The difference between MI and our IGCD-u setting is that in MI it is assumed that at each stage there are always instances from the previous stage's categories in the unlabeled images.

GM [56] only report results for the MI setting on CIFAR-100. For fairness, we re-train their model using the same hyper-parameters as ours and observe improved performance for their method compared to their paper. We also present additional results on CUB [53]. We include com-

Methods	Stage-0	Stage-1		Stage-2			Stage-3				Overall	
	All	New	S-0	New	S-1	S-0	New	S-2	S-1	S-0	$M_f \downarrow$	M_d
Supervised upper-bound	42.3	59.8	48.7	62.3	63.9	49.6	64.3	64.3	64.1	52.4	-	-
SimGCD [54]	42.3	8.4	5.1	12.6	2.3	0.1	11.8	4.2	1.2	0.0	42.2	2.5
SimGCD + iCaRL [39]	41.5	10.5	16.7	13.4	8.4	11.2	14.5	10.2	7.1	10.2	31.3	12.4
FRoST [42]	42.3	6.7	12.3	9.2	7.4	9.3	10.0	8.4	7.1	5.1	37.2	7.3
GM [56]	42.3	8.6	18.4	13.8	9.1	**15.5**	10.7	11.0	8.2	13.0	29.3	10.5
Ours	41.8	**12.7**	**22.1**	**15.4**	**10.7**	14.3	**16.1**	**11.2**	**9.1**	**13.1**	**28.7**	**14.1**

Table 2: Results on iNatIGCD in the IGCD-u setting (i.e. where labeled data is not provided during the incremental stages). Higher numbers are better, with the exception of M_f where lower is better.

parisons to the strong GCD baseline SimGCD [54], which cannot learn incrementally, and our incremental extension of their method. We observe that SimGCD [54] exhibits catastrophic forgetting, but adding iCaRL [39] can alleviate this issue. Our proposed method achieves better performance than GM [56], with the exception of the M_f performance on the smaller CIFAR-100 dataset.

Methods	CIFAR-100		CUB	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
SimGCD [54]	58.7	28.3	63.5	27.4
SimGCD + iCaRL [39]	9.4	29.4	10.7	28.3
GM [56]	**3.6**	30.6	6.8	26.7
Ours	4.0	**31.2**	**6.7**	**29.4**

Table 3: Results on the MI setting introduced in [56].

5.2.3 Ablations

Estimating the Number of Novel Categories. One of the unique challenges in GCD is the task of estimating the number of novel categories in the unlabeled data. [51] described a baseline that uses semi-supervised k-means to estimate it. GM [56] adopted this method to estimate the number at each incremental stage before training their classifier. Our proposed density selection method can also be seen as a way of estimating the number of categories based on the number of density peaks. In Tab. 4, we demonstrate that our method is capable of providing a more accurate estimate compared to GM, without requiring their expensive k-means step. Note, the ground truth class count for iNatGCD-l in Tab. 4 indicates the number of all categories including 'Old' and 'New', thus it is different from the counts in Fig. 2.

Impact of Selection Method for $\mathcal{R}$. Compared to the centroid selection in iCaRL [39], our method uses the concept of density peaks to select the examples to save in the replay buffer $\mathcal{R}$. In Tab. 5 we compare centroid selection and our proposed density-based selection. We observe that our method achieves superior performance.

Impact of the Size of $\mathcal{S}$. One of the key factors influencing the performance of our SNN classifier is the size of $\mathcal{S}$. In Tab. 6, we study the influence of this hyper-parameter by varying the number of examples per category ($N^{\mathcal{S}}/K^{\mathcal{S}}$).

CIFAR-100 MI	# Classes ($t=1$)	$t=2$	$t=3$
Ground truth	10	10	10
GM [56]	14	13	13
Ours	13	10	12
iNatIGCD-l	# Classes ($t=1$)	$t=2$	$t=3$
Ground truth	972	3,040	4,324
GM [56]	857	2,563	3,854
Ours	886	2,857	4,085

Table 4: Estimation of the number of novel categories in the unlabeled data at each stage.

Methods	CIFAR-100		iNatIGCD-l	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
Ours w/ iCaRL [39]	**3.4**	30.7	28.3	27.1
Ours	4.0	**31.2**	**27.6**	**28.4**

Table 5: Impact of the replay buffer selection method.

When only using one example per category, our model exhibits forgetting and inferior discovery performance. Increasing the number of examples per category increases the performance, until it starts to plateau after five examples per category. To reduce the cost of saving too many examples to $\mathcal{S}$, we set $N^{\mathcal{S}}/K^{\mathcal{S}}$ to five in our experiments.

$N^{\mathcal{S}}$ / $K^{\mathcal{S}}$	CIFAR-100		iNatIGCD-l	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
1	20.4	16.7	39.6	15.7
3	8.5	25.7	31.2	20.5
5	4.0	31.2	27.6	28.4
7	4.1	32.0	27.0	28.6
10	**3.1**	**33.2**	**25.1**	**29.5**

Table 6: Impact of the size of the support set $\mathcal{S}$.

6. Conclusion

We explored the problem of incremental generalized category discovery (IGCD). To do this, we constructed a new dataset, iNatIGCD, motivated by a real-world fine-grained visual discovery task and used it to benchmark the performance of recent category discovery methods. Through our experiments, we showed that our new approach which combines non-parametric categorization with a density-based

sample selection technique is superior to existing methods. While promising, our results show that the IGCD problem, especially in our large-scale fine-grained setting, remains challenging for existing methods. We hope that our work opens the door to future progress on this task.

References

[1] F Gregory Ashby and W Todd Maddox. Human category learning. *Annual review of psychology*, 2005. 1

[2] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, 2022. 2

[3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *ICCV*, 2021. 2, 5

[4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 2

[5] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 3, 4

[6] Kaidi Cao, Maria Brbić, and Jure Leskovec. Open-world semi-supervised learning. In *ICLR*, 2022. 1, 2, 13

[7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2

[8] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. The MIT Press, 2006. 2

[9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2, 5, 12

[10] Yuzhao Chen, Zonghuan Li, Zhiyuan Hu, and Nuno Vasconcelos. Taxonomic class incremental learning. *arXiv:2304.05547*, 2023. 3

[11] Florent Chiaroni, Jose Dolz, Ziko Imtiaz Masud, Amar Mitiche, and Ismail Ben Ayed. Mutual information-based generalized category discovery. *arXiv:2212.00334*, 2022. 2

[12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 6, 11

[13] Yixin Fei, Zhongkai Zhao, Siwei Yang, and Bingchen Zhao. Xcon: Learning with experts for fine-grained category discovery. In *BMVC*, 2022. 3

[14] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *ICCV*, 2021. 4, 13

[15] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *ICLR*, 2020. 2, 3, 6

[16] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *ICCV*, 2019. 1, 2, 3, 6

[17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 6

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6, 11

[19] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. Learning to cluster in order to transfer across domains and tasks. In *ICLR*, 2018. 2

[20] Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. Multi-class classification without multi-class labels. In *ICLR*, 2019. 2

[21] Junkai Huang, Chaowei Fang, Weikai Chen, Zhenhua Chai, Xiaolin Wei, Pengxu Wei, Liang Lin, and Guanbin Li. Trash to treasure: harvesting ood data with cross-modal matching for open-set semi-supervised learning. In *ICCV*, 2021. 2

[22] iNaturalist. iNaturalist. www.inaturalist.org, accessed 8 March 2023. 2, 3

[23] KJ Joseph, Sujoy Paul, Gaurav Aggarwal, Soma Biswas, Piyush Rai, Kai Han, and Vineeth N Balasubramanian. Novel class discovery without forgetting. In *ECCV*, 2022. 1, 3, 14

[24] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020. 5, 12

[25] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *PNAS*, 2017. 2

[26] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Workshop on 3D Representation and Recognition*, 2013. 2, 3

[27] Alex Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*, 2009. 2, 3

[28] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017. 2

[29] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 2015. 1

[30] Suichan Li, Bin Liu, Dongdong Chen, Qi Chu, Lu Yuan, and Nenghai Yu. Density-aware graph for deep semi-supervised visual recognition. In *CVPR*, 2020. 2, 6

[31] Zhiqiu Lin, Deepak Pathak, Yu-Xiong Wang, Deva Ramanan, and Shu Kong. Continual learning with evolving class ontologies. *NeurIPS*, 2022. 3

[32] Mingxuan Liu, Subhankar Roy, Zhun Zhong, Nicu Sebe, and Elisa Ricci. Large-scale pre-trained models are

surprisingly strong in incremental novel class discovery. *arXiv:2303.15975*, 2023. 1, 3

[33] Yu Liu and Tinne Tuytelaars. Residual tuning: Toward novel category discovery without labels. *Transactions on Neural Networks and Learning Systems*, 2022. 3

[34] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 13

[35] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv:1306.5151*, 2013. 2, 3

[36] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, 1989. 2

[37] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *ICPR*, 2006. 6

[38] Avital Oliver, Augustus Odena, Colin A Raffel, Ekin Dogus Cubuk, and Ian Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, 2018. 2

[39] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017. 1, 3, 4, 6, 7, 8, 12, 13

[40] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *ICLR*, 2021. 2

[41] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 1995. 3, 4

[42] Subhankar Roy, Mingxuan Liu, Zhun Zhong, Nicu Sebe, and Elisa Ricci. Class-incremental novel class discovery. In *ECCV*, 2022. 1, 3, 7, 8, 14

[43] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 3

[44] Kuniaki Saito, Donghyun Kim, and Kate Saenko. Openmatch: Open-set consistency regularization for semi-supervised learning with outliers. In *NeurIPS*, 2021. 2

[45] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boult. Toward open set recognition. *PAMI*, 2012. 1

[46] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020. 2

[47] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2

[48] Sebastian Thrun. Is learning the n-th thing any easier than learning the first? In *NeurIPS*, 1995. 1

[49] Colin Troisemaine, Vincent Lemaire, Stéphane Gosselin, Alexandre Reiffers-Masson, Joachim Flocon-Cholet, and Sandrine Vaton. Novel class discovery: an introduction and key concepts. *arXiv:2302.12028*, 2023. 1

[50] Grant Van Horn, Elijah Cole, Sara Beery, Kimberly Wilber, Serge Belongie, and Oisin Mac Aodha. Benchmarking representation learning for natural world image collections. In *CVPR*, 2021. 3, 4

[51] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *CVPR*, 2022. 1, 2, 3, 5, 6, 8, 13

[52] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: a good closed-set classifier is all you need? In *ICLR*, 2022. 3

[53] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Caltech-ucsd birds 200. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. 2, 3, 7

[54] Xin Wen, Bingchen Zhao, and Xiaojuan Qi. A simple parametric classification baseline for generalized category discovery. *arXiv:2211.11727*, 2022. 2, 4, 5, 7, 8, 12, 13

[55] Qing Yu, Daiki Ikami, Go Irie, and Kiyoharu Aizawa. Multi-task curriculum framework for open-set semi-supervised learning. In *ECCV*, 2020. 2

[56] Xinwei Zhang, Jianwen Jiang, Yutong Feng, Zhi-Fan Wu, Xibin Zhao, Hai Wan, Mingqian Tang, Rong Jin, and Yue Gao. Grow and merge: A unified framework for continuous categories discovery. In *NeurIPS*, 2022. 1, 2, 3, 4, 6, 7, 8, 11, 12, 13, 14

[57] Bingchen Zhao and Kai Han. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. In *NeurIPS*, 2021. 3, 6

[58] Da-Wei Zhou, Qi-Wei Wang, Zhi-Hong Qi, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Deep class-incremental learning: A survey. *arXiv:2302.03648*, 2023. 3

[59] Xiaojin Jerry Zhu. Semi-supervised learning literature survey. *University of Wisconsin-Madison Department of Computer Sciences*, 2005. 1, 2

A. Additional Results

A.1. Additional Ablations

Impact of the Size of $\mathcal{R}$. In addition to the ablation on the size of the support set $\mathcal{S}$ in the main paper, here we present results where we vary the size of the replay buffer $\mathcal{R}$. Similar to the ablation of the size of $\mathcal{S}$, we denote the number of images in $\mathcal{R}$ as $N^{\mathcal{R}}$ and the number of categories in $\mathcal{R}$ as $K^{\mathcal{R}}$. The results are presented in Tab. A1. The default value used by our method is illustrated by the gray row, i.e. the second row. We observe a similar trend as Table 6 in the main paper where increasing the number of examples increases the performance. However, the benefit starts to plateau and the larger number impacts training efficiency especially in the context of the large number of categories in our iNatIGCD dataset.

$N^{\mathcal{R}}$ / $K^{\mathcal{R}}$	CIFAR-100		iNatIGCD-l	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
1	25.1	18.6	38.2	15.8
3	4.6	29.2	28.3	27.1
5	4.0	31.2	28.1	26.6
7	4.4	34.2	27.1	27.9
10	**3.4**	**36.1**	**24.0**	**29.4**

Table A1: Impact of the size of the replay buffer $\mathcal{R}$.

Density Selection Parameters. There are three hyper-parameters to be set in our density selection step described in Sect. 4.3 in the main paper. They are, the number of nearest neighbors K for calculating the density, the number of neighbors K^d, and the threshold T used when removing redundant density peaks. We present the ablation study on these parameters in Tabs. A2 to A4.

In Tab. A2, we vary the value of K. We can see that generally the performance of different K values results in an inverse U shape. As K increases, the performance reaches the best value, and after the performance peak, the performance degrades as K increases further. When K is small, the compared neighborhood is small, thus the estimation of density peaks is not accurate. This results in many noisy peaks that can lead worse performance. When K increases, the density peak estimation is more accurate. However, considering the case where K equals to the number of images in the dataset, we would underestimate the density peaks as many clusters may be considered as one cluster. This would result in a performance decrease when K is higher than better choices.

In Tab. A3, we ablate the choice of K^d. This hyper-parameter is used to remove redundant density peaks. When the value of K^d is low, it would remove fewer density peaks as two peaks' K^d neighbor set are less likely to overlap. When the value is high, it would remove more density peaks. From the results, we can see that setting K^d to a higher number than K (which equals 10 in Tab. A3) can result in better performance.

K	CIFAR-100		iNatIGCD-l	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
5	4.8	27.1	28.9	27.1
10	4.3	28.9	27.9	26.9
15	4.5	**29.7**	26.1	**27.8**
20	4.1	28.1	**25.9**	27.6
40	**3.8**	26.4	27.2	26.0

Table A2: Ablation of K used in density-based selection.

K^d	CIFAR-100		iNatIGCD-l	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
10	4.1	20.3	28.0	18.9
15	4.8	23.2	27.5	23.4
20	4.3	28.9	27.9	26.9
30	**3.8**	**29.5**	27.1	**27.9**
40	4.5	28.1	**26.9**	27.2

Table A3: Ablation of K^d used in density-based selection.

Another hyper-parameter in density selection is the threshold T. Similar to the selection of K^d, a higher threshold will remove fewer density peaks, while a lower threshold will remove more. The results are presented in Tab. A4. Our choice from the main paper is shaded in gray. This obtains a good balance between forgetting and discovery.

T	CIFAR-100		iNatIGCD-l	
	$M_f \downarrow$	$M_d \uparrow$	$M_f \downarrow$	$M_d \uparrow$
0.2	5.1	28.7	29.2	**27.0**
0.4	4.8	**29.3**	28.9	26.8
0.6	4.3	28.9	27.9	26.9
0.8	**4.0**	27.2	**27.1**	26.0

Table A4: Ablation of T used in density-based selection.

A.2. Vision Transformer Results

In the main paper, our results are presented using a ResNet-18 backbone. In Tab. A5 we present results using a vision transformer [12] Base 16 model. We observe that the overall trend remains the same. The vision transformer results in an increase in performance compared to the ResNet-18 [18] that we use in the main paper, but at the cost of being much slower to train.

A.3. CIFAR-100 IGCD-l and IGCD-u Results

In Tabs. A6 and A7, we re-purpose the CIFAR-100 dataset for the IGCD-l and IGCD-u settings, using the same category split as [56], and present the results comparing our method with previous SoTA baselines. We can see that our method, when compare with others, still achieves a competitive performance. In terms of the overall performance

Methods	Stage-0	Stage-1			Stage-2				Stage-3					Overall	
	All	All	Old	New	All	Old	New	S-0	All	Old	New	S-1	S-0	M_f	M_d
SimGCD + iCaRL [39]	45.6	25.2	33.6	14.1	24.7	**38.0**	16.9	9.7	24.0	37.0	17.8	15.3	9.7	35.9	26.9
GM [56]	**46.0**	19.5	28.1	9.4	20.9	29.6	11.2	16.2	18.3	27.1	16.8	17.0	14.3	**29.1**	24.8
Ours	45.6	**27.0**	**36.3**	**16.7**	**26.8**	37.2	**23.4**	**18.5**	**25.0**	**38.1**	**19.7**	**18.2**	**17.1**	29.4	**30.2**

Table A5: Vision transformer-backbone results on iNatIGCD in the IGCD-l setting (i.e. where labels are available at the end of each stage). Higher numbers are better, with the exception of M_f, where lower is better.

```
# feats: feature vectors of the unlabeled set (NxC).
# supps: existing support samples, normed (SxC)
# K: the number of neighbours to use for density
# Kd: the number of neighbours to use for IoU
# thres: threshold for IoU and NMS
# returns selected_peaks

# get k nearest neighbors
normed_feats = norm(feats)
sims = normed_feats @ normed_feats.t()
knn_indice = topk(sims, K)

# get density peaks
density = sims[knn_indice].sum() / K
nn_density = density[knn_indice]
peaks_ind = [i where density[i] > nn_density[i]]

# use iou on nns to dedup
knn_ind = topk(sims, k=Kd)
peak_nns = knn_ind[peaks_ind]
selected_peaks = [peaks_ind[0]]
for peak, nn in zip(peaks_ind[1:], peak_nns):
    flag = True
    for other_nn in peak_nns.remove(nn):
        if IoU(nn, other_nn) > thres:
            flag = False
    if flag:
        selected_peaks.append(peak)
```

Listing 1: Pseudo-code of our density selection step.

metrics $\mathcal{M}_f$ and $\mathcal{M}_d$, our method achieves the best results for both CIFAR-100-IGCD-l and CIFAR-100-IGCD-u.

A.4. Results on Static GCD Datasets

In Tab. A8 we present results on the static GCD benchmarks, i.e. without any incremental stages. For the adaptation of our method to the static GCD scenario, we remove the incremental update of $\mathcal{S}$ and $\mathcal{R}$. We can see that our proposed method still performs on par with previous state-of-the-art methods on static GCD benchmarks despite primarily being designed for the incremental setting.

A.5. Performance of Methods Over Stages

B. Implementation Details

B.1. Training Losses

We describe the representation learning losses $\mathcal{L}_{\text{rep}}$ below which follows a contrastive learning framework. Formally, given two views (i.e. augmentations) $\hat{\boldsymbol{x}}_i$ and $\tilde{\boldsymbol{x}}_i$ of an input image $\boldsymbol{x}_i$ in a mini-batch $\mathcal{B}$, the self-supervised component of the contrastive loss can be written as

$$\mathcal{L}_{\text{SelfCon}} = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} -\log \frac{\exp\left(\hat{\boldsymbol{z}}_i^\top \tilde{\boldsymbol{z}}_i' / \tau_u\right)}{\sum_i^{i \neq n} \exp\left(\hat{\boldsymbol{z}}_i^\top \tilde{\boldsymbol{z}}_n' / \tau_u\right)}, \quad (10)$$

where the embedding $\boldsymbol{z} = m(f(\boldsymbol{x}))$ is a projected feature from a MLP projector m as in [9, 54], and τ_u is a temperature parameter. The supervised contrastive loss [24] is similar with the difference being that the positive samples are matched with their ground truth labels,

$$\mathcal{L}_{\text{SupCon}} = \frac{1}{|\mathcal{B}^l|} \sum_{i \in \mathcal{B}^l} \frac{1}{|\mathcal{M}_i|} \sum_{q \in \mathcal{M}_i} -\log \frac{\exp\left(\hat{\boldsymbol{z}}_i^\top \tilde{\boldsymbol{z}}_q / \tau_c\right)}{\sum_n^{n \neq i} \exp\left(\hat{\boldsymbol{z}}_i^\top \tilde{\boldsymbol{z}}_n / \tau_c\right)}, \quad (11)$$

where $\mathcal{M}_i$ indexes all other images in the same batch that have the same label as $\boldsymbol{x}_i$, and τ_c is the temperature parameter for the supervised contrastive loss.

For the fully supervised upper bound discussed in Tabs. 1 and 2 in the main paper, we leverage both contrastive learning losses $\mathcal{L}_{\text{SelfCon}}$ and $\mathcal{L}_{\text{SupCon}}$ for representation learning. For the classifier, we simply adopt the cross-entropy loss $\mathcal{L}_{\text{ce}}$ on all examples. We concatenate all datasets $\mathcal{D}_{lab}^t, \mathcal{D}_{unlab}^t$ from every time step t, and provide the ground truth category labels for the unlabelled datasets $\mathcal{D}_{unlab}^t$. The resulting concatenated dataset is used to train the fully supervised baseline. So in this setting, the model will have labels for all the categories and no forgetting occurs. This can be considered as a very strong upper bound on the performance for our IGCD setting.

B.2. Density Selection

We present the pseudo-code for our proposed density selection process in Listing 1. We use ℓ_2-normed features to calculate the density of the examples, and then filter out the density peaks by comparing the density of one node to its k nearest neighbours. Then we compare each pair of density peaks and remove peaks that have nearest neighbours set to overlap with other peaks larger than a threshold.

B.3. Evaluation Metrics

The main results in each of the tables and figures are reported using the clustering accuracy (ACC). At each evaluation stage, given the ground truth $\boldsymbol{y}$ and a predicted label $\hat{\boldsymbol{y}}$, the ACC is calculated as $\text{ACC} = \frac{1}{M} \sum_{i=1}^{M} \mathbb{1}(\boldsymbol{y}_i = p(\hat{\boldsymbol{y}}_i))$.

Methods	Stage-0	Stage-1			Stage-2				Stage-3					Overall	
	All	All	Old	New	All	Old	New	S-0	All	Old	New	S-1	S-0	M_f	M_d
Supervised upper-bound	79.2	82.5	80.3	85.3	81.0	82.4	80.0	79.2	81.3	82.3	80.2	81.7	80.2	-	-
SimGCD + iCaRL [39]	77.8	**78.6**	**79.6**	76.7	77.9	78.6	74.5	63.4	75.6	**78.9**	70.8	60.2	50.1	27.7	67.8
GM [56]	79.0	78.3	78.8	77.3	**78.2**	**78.9**	74.7	**66.3**	**76.7**	77.8	71.2	63.2	49.8	29.2	68.3
Ours	77.8	78.5	79.0	78.2	78.0	78.7	**75.0**	65.3	76.0	78.8	**72.1**	**64.6**	**51.3**	**26.5**	**69.2**

Table A6: Results on CIFAR-100 in the IGCD-l setting (i.e. where labels are available at the end of each stage). Higher numbers are better, with the exception of M_f, where lower is better.

Methods	Stage-0	Stage-1		Stage-2			Stage-3				Overall	
	All	New	S-0	New	S-1	S-0	New	S-2	S-1	S-0	M_f	M_d
Supervised upper-bound	79.2	78.9	81.2	80.2	79.9	81.4	85.3	82.1	80.1	79.0	-	-
SimGCD + iCaRL [39]	77.8	62.4	43.2	66.7	40.8	38.7	**68.2**	41.3	36.2	35.3	42.5	50.1
GM [56]	79.0	58.9	45.3	63.4	46.8	**43.2**	60.0	**42.5**	45.6	40.2	38.8	53.4
Ours	77.8	**64.4**	**46.9**	**67.8**	**47.0**	42.1	68.0	40.0	**46.8**	**41.5**	**36.3**	**55.9**

Table A7: Results on CIFAR-100 in the IGCD-u setting (i.e. where labeled data is not provided during the incremental stages). Higher numbers are better, with the exception of M_f, where lower is better.

Methods	ImageNet-100			SCars		
	All	Old	New	All	Old	New
GCD [51]	74.1	89.8	66.3	39.0	57.6	29.9
ORCA [6]	73.5	**92.6**	63.9	23.5	50.1	10.7
SimGCD [54]	82.4	90.7	78.3	46.8	**64.9**	38.0
Ours	**83.0**	89.5	**79.1**	**47.2**	63.8	**38.7**

Table A8: Results on static GCD benchmarks.

Here, $M = |\mathcal{D}^u|$ and p is the optimal permutation that matches the predicted cluster assignments to the ground truth category labels. Under the IGCD-l setting, we report the ACC for all the categories in Stage-0 as all categories have labeled training examples. For later stages, we report the 'All', 'Old', and 'New' performance which correspond to all categories, the labeled categories, and the unlabeled categories at each stage. Additionally, we use S-t to denote the categories that appear in a previous stage t but are not presented in the current stage to measure the forgetting of the model. Finally, as noted in the main paper we also use the summary metrics defined in [56] to measure overall discovery and forgetting performance, $\mathcal{M}_d$ and $\mathcal{M}_f$.

B.4. Experimental Settings

Our models are trained using a stochastic gradient descent optimizer with an initial learning rate of 0.1, momentum of 0.9, and weight decay of $1e-4$. The learning rate is decayed following a cosine schedule [34]. We use a batch size of 128 during training for all datasets, where 64 images are sampled from $\mathcal{D}^t_{lab}$ and 64 images are sampled from $\mathcal{D}^t_{unlab}$. The balancing factor λ_{rep} is set to 0.35, ϵ is set to 2.0, K is set to 10, K^d is set to 20, and the threshold T is set to 0.6. We adopt the same set of augmentations used in [14] for contrastive learning.

C. iNatIGCD Dataset

In Tab. A9, we present a comparison of our dataset split to common benchmarks used by previous papers. Our new iNatIGCD benchmark has more categories and images per stage and thus can be used to better evaluate the performance of generalized category discovery methods. For completeness, in Fig. A1 we visually illustrate the difference between the IGCD-l and IGCD-u settings.

In Fig. A2, we depict visual examples from our proposed dataset to visualize the difference in the appearance of the same species across different locations. We observe that across different locations, the appearance or the background of species may change dramatically which provides an additional challenge to the learner.

Dataset	Used In	# Classes / Stage	# Stages	# Avg. Images / Stage	# Avg. Classes / Stage
ImageNet-1k	NCDwF [23]	882 / 30	2	39k	30
TinyImageNet	class-iNCD [42]	180 / 20	2	20k	20
CIFAR100-MI	GM [56]	70 / 10 / 10 / 10	4	5k	10
iNatIGCD	Ours	3,888 / 972 / 3,040 / 4,324	4	25k	2.7k

Table A9: The statistics of our dataset compared to previous benchmarks that focus on generalized category discovery under the incremental setting. At Stage-1, 2, and 3, iNatIGCD contains 390, 2,625, and 1,492 novel classes respectively.

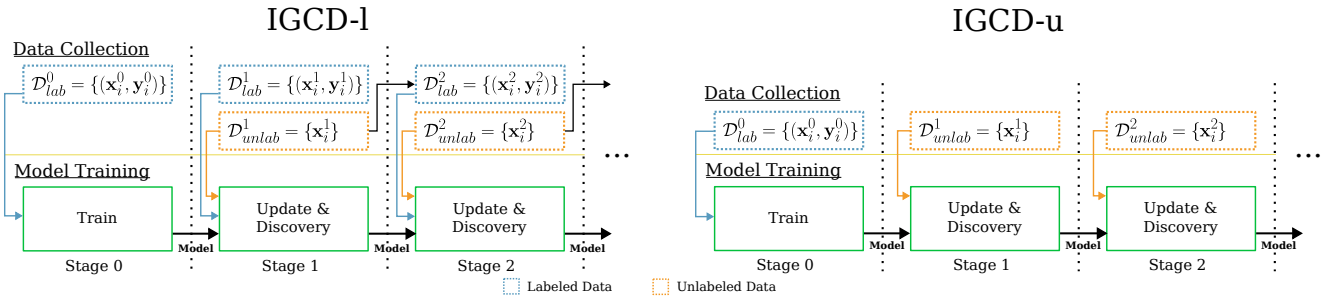

Figure A1: The two incremental learning settings used in our work. In the IGCD-l setting, at each stage we have one unlabeled set and one labeled set. The unlabeled set is annotated at the end of each stage and then used as the labeled set for the next stage. In the IGCD-u setting, only the unlabeled set is provided at each incremental stage, i.e. after the initial label data is given in Stage-0, no additional labels are provided.

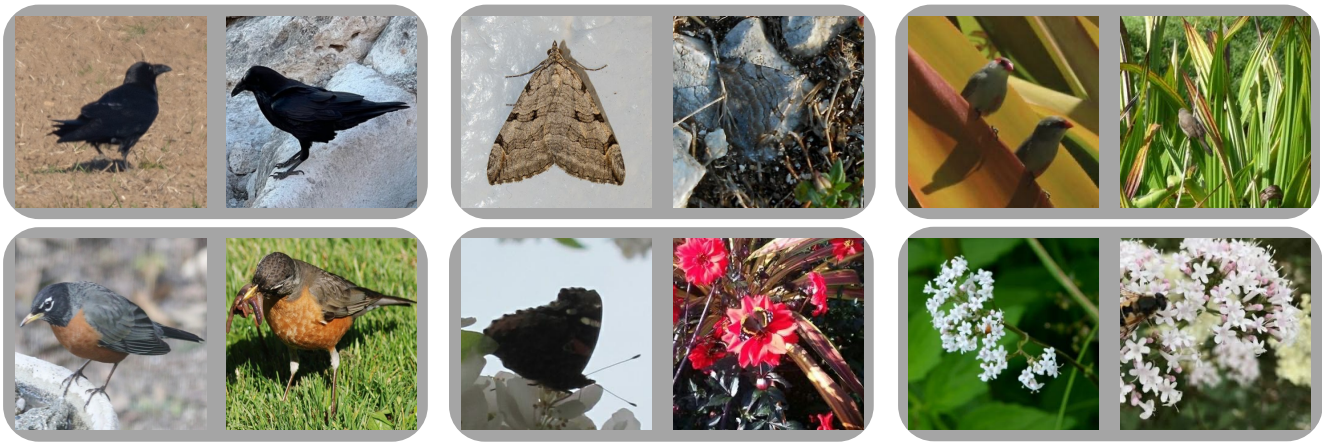

Figure A2: Example images from our iNatIGCD dataset, where we illustrate the difference between the same species across different continents. Each box contains two instances of the same species from different locations. We can observe that the background or appearance can change dramatically across different locations.