

Labeled Data Selection for Category Discovery

Bingchen Zhao¹ Nico Lang² Serge Belongie² Oisín Mac Aodha¹

¹University of Edinburgh ²University of Copenhagen

Abstract. Visual category discovery methods aim to find novel categories in unlabeled visual data. At training time, a set of labeled and unlabeled images are provided, where the labels correspond to the categories present in the images. The labeled data provides guidance during training by indicating what types of visual properties and features are relevant for performing discovery in the unlabeled data. As a result, changing the categories present in the labeled set can have a large impact on what is ultimately discovered in the unlabeled set. Despite its importance, the impact of labeled data selection has not been explored in the category discovery literature to date. We show that changing the labeled data does indeed significantly impact discovery performance. Motivated by this, we propose two new approaches for automatically selecting the most suitable labeled data based on the similarity between the labeled and unlabeled data. Our observation is that, unlike in conventional supervised transfer learning, the most informative labeled data is neither too similar nor too dissimilar, to the unlabeled categories. Our resulting approaches obtain state-of-the-art discovery performance across a range of methods and challenging fine-grained benchmark datasets.

Keywords: Data Selection · Category Discovery · GCD

1 Introduction

Given a set of labeled and unlabeled images, the goal of category discovery methods is to automatically group images in the unlabeled data into categories. The problem has been explored under several different guises, including settings where the unlabeled data only contains images from novel categories [20], where the unlabeled data can contain both previously seen *and* novel categories [45], and also in the online learning setting [52, 54]. What differentiates discovery from conventional unsupervised clustering [6, 21] is the addition of the labeled data.

On first appearance, the availability of labeled data may appear to only be a minor difference from traditional unsupervised clustering. However, the labeled data plays an important role for discovery algorithms. It provides guidance during training regarding the types of visual features that are deemed important and relevant for defining categories in the unlabeled data. For example, suppose one was provided with a dataset that contains images captured in nature. If the labeled data only contained images of different species of birds, this strongly limits the discovery method to only be able to find ‘bird-like’ concepts in the




Fig. 1: Discovery benefits from selecting a subset of the labeled data. Category discovery methods rely on labeled source data to provide context for discovering visual concepts in unlabeled target data. In the related task of supervised categorization via transfer learning [9], performance on the target data is maximized by selecting the most related concepts from the source data to pre-train the model (orange line). We show that discovery methods behave very differently (pink line). Learning representations from source data that is too similar to the target actually hinders discovery. Instead, it is preferable to select data that is neither too similar nor too dissimilar. The arrows \downarrow \downarrow \downarrow (right) indicate the locations of the subsets (left) with the same color.

unlabeled data. Similarly, if the labeled data instead only contained images of plant species, it would restrict the method to discover ‘plant-like’ concepts.

As a result, the role and selection of labeled data is an important question in category discovery. However, until now, this problem has been overlooked in the discovery literature, which has instead focused on developing new methods that can be applied to off-the-shelf datasets with already provided labeled and unlabeled splits. In the context of supervised classification, [9] showed that the choice of labeled pre-training source data can have a large impact on downstream fine-tuned target performance. Motivated by this, they proposed an approach for selecting a source training set that is most similar to a target domain of interest by measuring the similarity between the source and target domains. Their observation is that the ideal source dataset for supervised classification is the one that is most similar to the target domain. In this work, we show that this is *not* the case for category discovery.

In discovery, we demonstrate that the similarity between the labeled categories and the unlabeled data exposes a surprising, and at first perhaps counter-intuitive, relationship. Intuitively, if the categories in the labeled data are too *dissimilar* to the ones in the unlabeled data, discovery fails, *i.e.* labeled images of insects do not provide useful cues for discovering species of birds. However, if the labeled categories are too *similar* to the unlabeled ones this also dramatically hurts discovery performance, *i.e.* if the labeled and unlabeled datasets both contain very visually similar, but distinct, fine-grained bird categories, discovery methods are unable to separate the two sets of similar categories. As a result, there appears to be a ‘sweet spot’ in the middle of both extremes, whereby the labeled images should be related enough to provide useful context, but not too similar as to be overly difficult to separate. We show that exploiting this observation to select a subset of labeled categories for training category discovery methods results in significantly enhanced discovery performance (see Fig. 1).

We introduce a new labeled dataset selection framing of the generalized category discovery problem. At training time we are presented with an unlabeled target dataset in which we want to discover novel categories. In addition, we are given a labeled source dataset where the goal is to automatically select a subset from it to train the discovery algorithm. Importantly, this labeled dataset can potentially contain images from diverse categories, many of which may not be relevant to the unlabeled target. We propose two new unsupervised selection methods that assign weight to the labeled data during the training process so that labeled data that is too similar or too dissimilar to the target data is down-weighted and thus has less influence on the learned representation.

We make the following contributions: (i) We show that category discovery performance highly depends on the specific labeled data used during training. (ii) We propose two new methods for automatically selecting a more effective labeled training dataset to improve discovery performance. (iii) We show that our approaches generalize across multiple discovery methods and result in state-of-the-art performance on several challenging fine-grained discovery benchmarks.

2 Related work

2.1 Category discovery

The task of category discovery centers around the development of algorithms that aim to find novel categories in unlabeled data. In addition to unlabeled data, labeled data is also provided at training time which gives information indicating the types of visual concepts that are to be discovered. Several different variants of the core category discovery problem have been proposed in the literature.

In novel category discovery (NCD), the unlabeled data is assumed to not contain any images from the categories in the labeled set [15, 18–20, 53, 56]. This is a overly simplifying assumption, and in the more realistic generalized category discovery (GCD) setting the unlabeled data can contain images from both previously seen and novel categories [5, 14, 33, 45, 47, 49, 51]. Both of these discovery paradigms operate under a closed set assumption, *i.e.* there are no novel categories in the test set that are not in the labeled or unlabeled sets. On-the-fly category discovery (OCD) attempts to handle this specific challenge [12, 34]. Finally, there is a recent set of methods that explore discovery in the online learning setting [24, 52, 54]. Here, the goal is to not only discover new concepts but also not to forget previously seen ones.

The above methods have primarily focused on methodological improvements and have neglected to investigate the role of training data. Theoretical works on category discovery have also emerged of late. [7] examined the bound for category discovery when the number of examples is limited and [38] and [39] investigated how labeled data helps the discovery of novel categories from a spectral analysis perspective. In this work, we explore the impact of the labeled data in the generalized discovery setting. We show that the choice of which labeled data to use can significantly impact downstream discovery performance. Based on this, we outline a simple and practical solution for selecting labeled data

for category discovery. Our results, across a range of recent methods, illustrate that state-of-the-art discovery performance can be obtained by simply training on selected labeled data.

2.2 Labeled data selection

The conventional wisdom in supervised categorization is that more data, coupled with larger models, results in improved performance [10, 37]. However, not all data is equally useful and commonly used supervised training datasets in vision have required large amounts of manual effort to collect and curate [26, 35, 42, 44]. Existing work has shown that both supervised [13, 50] and self-supervised [32] approaches can benefit from semi-automated methods for data filtering when constructing training datasets. There is also a set of works that aim to select a subset of a larger labeled training set, where the goal is to reduce the amount of training time needed [22, 27, 30].

In the context of supervised transfer learning, [2] demonstrated that transfer performance is dependent on the similarity between the source and target domains. [9] proposed a method for selecting a subset of labeled images from a large source collection such that they would be a useful pre-training source for a model that is fine-tuned on a smaller target dataset of interest. Their approach estimates domain similarity between class prototypes from the source and target datasets using the Earth Mover’s Distance and then greedily selects images from the categories in the source that are most similar to the target. There are two issues associated with applying their method to the discovery problem. First, we do not have labels for the images in the unlabeled set, and as we will demonstrate later, selecting the most similar categories between the source and the target actually hurts discovery performance.

Also partially related to our work are methods that have explored the use of different granularities of supervision during training. [8] showed that the precise label hierarchy used during supervised pre-training can have a large impact on object localization performance. Specifically, they showed in the context of hierarchically labeled data that using labels that are too coarse (*e.g.* using a generic label such as ‘bird’) or too fine (*e.g.* using a specific bird species as the label) performs worse than intermediate supervision. The impact of label granularity has also been explored in tasks such as action [36] and attribute-based recognition [17]. Works that study the transferability of pre-trained models also exist [3, 4, 31, 41], but they cannot be used for category discovery as they are aiming to select models, not data, from a large model zoo. Finally, some works have shown that features from networks trained using self-supervision on ImageNet [35] are effective for discovery [53, 54]. However, to the best of our knowledge, the impact of choosing different labeled data sources for category discovery has not been investigated.

3 Method

3.1 Problem statement

In generalized category discovery we are provided with a source labeled dataset $\mathcal{D}^l = \{(\mathbf{x}_i, y_i)\}_{i=1}^N \in \mathcal{X} \times \mathcal{Y}_l$ and a target unlabeled dataset $\mathcal{D}^u = \{(\mathbf{x}_j, y_j)\}_{j=1}^M \in$

$\mathcal{X} \times \mathcal{Y}_u$. In our case, \mathbf{x}_i is an image and y_i is its corresponding discrete category label. The labels y_j for the unlabeled dataset are *not* available during training. The images in the unlabeled dataset can contain novel categories that are not present in the labeled set, $\mathcal{Y}_u \setminus \mathcal{Y}_l \neq \emptyset$. We also do not know the number of possible categories present in the unlabeled set $|\mathcal{Y}_u|$. Importantly, in the generalized setting [45], there can also be images in the unlabeled set from categories that are also present in the labeled set, $\mathcal{Y}_l \subset \mathcal{Y}_u$. The task is to develop a feature extractor $f(\cdot)$ that, when combined with a classifier $g(\cdot)$, correctly assigns category labels to the instances in the unlabeled dataset, *i.e.* $\mathbf{p}_i = g(f(\mathbf{x}_i))$, where \mathbf{p}_i is a distribution over category labels from both the unlabeled and labeled sets. A successful discovery method correctly classifies unlabeled images as being either categories from the labeled set, or groups them into novel categories.

A growing number of methods have been proposed to address the discovery problem [5, 33, 45, 49]. Instead of developing a novel discovery method, we ask an orthogonal question. Specifically, we ask: “*What impact does the choice of the labeled dataset \mathcal{D}^l have on downstream discovery performance?*” The motivation for this question is that in category discovery, not only do we not know the number of novel categories in the unlabeled data in advance [55], but we also may not even know what ‘types’ of categories are present in the unlabeled data. Specifically, the scenario we are considering in this work is one where we are given a large source collection of labeled data and we have to automatically select a subset of it to supervise a model for category discovery on a specific unlabeled target dataset. To address this, we propose two new approaches for automatically selecting a labeled training dataset to significantly improves discovery performance, all without altering the underlying discovery algorithm used.

3.2 Category discovery

For our main experiments, we make use of the recent SimGCD [49] discovery approach. We chose this method because it is conceptually simple and obtains state-of-the-art performance across a range of challenging datasets. We also later demonstrate that our observations hold for other discovery methods (*i.e.* GCD [45], XCon [14], and μ GCD [47]). In this section, we briefly describe the key components of SimGCD, which consists of two main parts: (i) representation learning and (ii) classifier learning.

Representation learning. To learn a discriminative representation that can separate and discover novel categories, SimGCD employs a pre-trained feature extractor $f(\cdot)$ and fine-tunes it with contrastive learning losses. Given two random augmentations $\hat{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}_i$ of an image \mathbf{x}_i in a training mini-batch B that contains both labeled and unlabeled data, the *self-supervised* contrastive loss can be computed as:

$$\mathcal{L}_{\text{rep}}^u = \frac{1}{|B|} \sum_{\mathbf{x}_i \in B} -\log \frac{\exp(\hat{\mathbf{z}}_i^\top \tilde{\mathbf{z}}_i / \tau_u)}{\sum_{\mathbf{x}_j \in B} \exp(\hat{\mathbf{z}}_j^\top \tilde{\mathbf{z}}_i / \tau_u)}, \quad (1)$$

where $\mathbf{z}_i = m(f(\mathbf{x}_i))$ is a projected feature for contrastive learning, $m(\cdot)$ is the projection head, and τ_u is a temperature value. In addition to this self-supervised

contrastive loss, SimGCD also makes use of a *supervised* contrastive loss to learn a discriminative representation from the labeled data:

$$\mathcal{L}_{\text{rep}}^s = \frac{1}{|B^l|} \sum_{\mathbf{x}_i \in B^l} \omega_i \frac{1}{|\mathcal{N}_i|} \sum_{p \in \mathcal{N}_i} -\log \frac{\exp(\hat{\mathbf{z}}_i^\top \tilde{\mathbf{z}}_p / \tau_s)}{\sum_{n \neq i} \exp(\hat{\mathbf{z}}_i^\top \tilde{\mathbf{z}}_n / \tau_s)}, \quad (2)$$

where \mathcal{N}_i is the indexes of all images in the mini-batch that have the same label as \mathbf{x}_i , τ_s is an additional temperature value, and ω_i is a weight factor for the contribution of \mathbf{x}_i to the learning process. We set the weight to 1 by default, and later discuss the impact of changing it to ‘select’ different data subsets. The representation learning training objective is the combined loss: $\mathcal{L}_{\text{rep}} = (1 - \lambda)\mathcal{L}_{\text{rep}}^u + \lambda\mathcal{L}_{\text{rep}}^s$, where λ is a weighting hyperparameter.

Classifier learning. SimGCD employs a parametric classifier to assign labels to images. The classifier is trained in a self-distillation fashion where the number of categories $K = |\mathcal{Y}_u|$ is given *a-prior*, or can be estimated using off-the-shelf methods, and SimGCD randomly initializes a set of prototypes for each category $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_K\}$. In the training process, a soft label $\hat{\mathbf{p}}_i$ is calculated by performing softmax classification using the prototypes:

$$\hat{\mathbf{p}}_i^k = \frac{\exp\left(\frac{1}{\tau_s}(\hat{\mathbf{h}}_i / \|\hat{\mathbf{h}}_i\|_2)^\top (\mathbf{c}_k / \|\mathbf{c}_k\|_2)\right)}{\sum_j \exp\left(\frac{1}{\tau_s}(\hat{\mathbf{h}}_i / \|\hat{\mathbf{h}}_i\|_2)^\top (\mathbf{c}_j / \|\mathbf{c}_j\|_2)\right)}, \quad (3)$$

where $\hat{\mathbf{h}}_i = f(\hat{\mathbf{x}}_i)$ is the representation of $\hat{\mathbf{x}}_i$ directly from the backbone $f(\cdot)$. In addition, a soft label vector $\tilde{\mathbf{p}}_i$ is also generated for $\tilde{\mathbf{x}}_i$. A cross-entropy and self-distillation with entropy-regularization loss are used to train this classifier:

$$\mathcal{L}_{\text{cls}}^l = \frac{1}{|B^l|} \sum_{\mathbf{x}_i \in B^l} \omega_i \mathcal{L}_{\text{ce}}(\mathbf{y}_i, \hat{\mathbf{p}}_i), \quad (4)$$

$$\mathcal{L}_{\text{cls}}^u = \frac{1}{|B|} \sum_{\mathbf{x}_i \in B} \mathcal{L}_{\text{ce}}(\tilde{\mathbf{p}}_i, \hat{\mathbf{p}}_i) - \epsilon H(\bar{\mathbf{p}}), \quad (5)$$

where \mathbf{y}_i is the ground truth label for the labeled example \mathbf{x}_i , \mathcal{L}_{ce} is the cross-entropy loss, ω_i is the weight factor defaulted to 1, and $H(\bar{\mathbf{p}}) = -\sum \bar{\mathbf{p}} \log \bar{\mathbf{p}}$ regularizes the mean prediction $\bar{\mathbf{p}} = \frac{1}{2|B|} \sum_{\mathbf{x}_i \in B} (\hat{\mathbf{p}}_i + \tilde{\mathbf{p}}_i)$ in a mini-batch. The final classifier loss is defined as $\mathcal{L}_{\text{cls}} = (1 - \lambda)\mathcal{L}_{\text{cls}}^u + \lambda\mathcal{L}_{\text{cls}}^l$.

The final overall SimGCD training loss is the combination of the representation learning and classifier learning losses:

$$\mathcal{L} = \mathcal{L}_{\text{rep}} + \mathcal{L}_{\text{cls}}. \quad (6)$$

3.3 Data selection

Our central hypothesis is that not all labeled data is equally useful for performing category discovery. To address this, we aim to ‘select’ a subset of the entire

labeled source dataset to use for training the discovery model (*e.g.* SimGCD). We perform data selection by generating instance-specific weights for each data point in the *labeled* source dataset. Specifically, the selection process aims to generate a weight $\omega_i \in [0, +\infty)$ for each data point \mathbf{x}_i in the labeled dataset which is used to weight some components of the training loss in Eqs. (2) and (4). Thus, data points with weight 0 are effectively discarded during training, and non-zero weighted data points influence training based on the magnitude of their weight. Similarly, using a weight of 1 for all data points is equivalent to running off-the-shelf discovery with no weighting. Next, we review a related supervised selection approach proposed for improving transfer learning and then introduce two new unsupervised selection methods.

Supervised data selection. In the context of supervised classification, [9] proposed a method for dataset selection where the goal was to improve transfer learning performance. Given a large labeled source dataset \mathcal{D}^s and a *labeled* target dataset \mathcal{D}^t , their method automatically selects a subset of the source data for pre-training an image classifier. The intuition is that if they can identify a subset of the source data that is most similar to the target dataset, it will provide a better pre-training signal for a model that can then be fine-tuned on the target dataset. They propose a method that uses the Earth Mover’s Distance (EMD) to measure the similarity between the source and target. As the number of images in the source dataset could be very large, they start with a simplifying assumption that each category can be represented by a single per-category mean feature vector, which we denote as:

$$\bar{\mathbf{h}}_c = \frac{1}{N_c} \sum_{\substack{(\mathbf{x}, y) \in \mathcal{D} \\ y=c}} f(\mathbf{x}), \quad (7)$$

where N_c is the number of instances of category c in a dataset \mathcal{D} . This mean category vector computation is only possible because both the source and target datasets are fully labeled. In the more challenging category discovery setting, which we will discuss later, the target data is unlabeled.

Given the per-category mean feature vectors they compute the EMD between two datasets as:

$$\text{EMD}(\mathcal{D}^s, \mathcal{D}^t) = \frac{\sum_{i=1}^N \sum_{j=1}^M k_{i,j} d_{i,j}}{\sum_{i=1}^N \sum_{j=1}^M k_{i,j}}, \quad (8)$$

where the distance between pairs of mean vectors is $d_{i,j} = \|\bar{\mathbf{h}}_i^s - \bar{\mathbf{h}}_j^t\|_2$, $k_{i,j}$ is the optimal flow that minimizes the overall cost, and N and M are the total number of categories in the source and target datasets. From this, they define a ‘domain similarity’ score between two datasets which is represented as:

$$\text{SIM}(\mathcal{D}^s, \mathcal{D}^t) = \exp(-\gamma \text{EMD}(\mathcal{D}^s, \mathcal{D}^t)), \quad (9)$$

where γ is a scaling hyperparameter.

To select a subset of the source dataset for pre-training they perform a greedy selection process whereby they incrementally select the categories from




Fig. 2: (Left) The beta distribution is used to generate weights for the labeled data. (Right) Overview of our labeled data selection process for category discovery. The labeled data is weighted based on the distance to the unlabeled data. The weight is used to change the influence of labeled categories during training. The green and orange arrows on the left panel denote the distance between labeled to unlabeled data for two instances illustrated using same color in the first step on the right panel.

the source that are most similar to the target. As this comparison is performed on the per-category mean feature vectors, to construct the final updated source dataset they select each of the images from the source that belong to the selected categories. Their approach can be viewed as assigning a weight of in $\{0, 1\}$ to each data point in the labeled dataset based on the category it belongs to.

Unsupervised data selection. In later category discovery experiments (Tab. 1) we demonstrate that the most informative labeled data to use is neither the most similar nor the most dissimilar to the *unlabeled* target. Instead, for discovery, it is more effective to select labeled data that shares high-level features with the target data such that learned features are helpful for discovery, and at the same time not too similar to cause confusion. If they were too similar, the model would likely incorrectly classify the novel categories in the unlabeled target as previously seen categories from the labeled dataset.

There are two main issues associated with applying the previously described selection approach from [9] to the problem of category discovery: (i) we do not have labels for the images in the unlabeled target data and (ii) the most similar categories from the source are not the most useful for category discovery (see Fig. 1). Next, we will outline two new conceptually simple unsupervised selection methods that address both of these challenges, where selection is performed by assigning generated weights (either hard or soft) to the labeled data.

Binning. One approach for data selection is to simply ‘bin’ the labeled dataset into several equal-sized subsets based on the similarity to the unlabeled target, and then select data from the bins that are not too similar or too dissimilar to the target. However, the difficulty is that without labels we cannot use the approach from [9] to determine similarity, thus we cannot determine how many bins to use or which bins to discard. To address this, we design an unsupervised selection method that first filters out unrelated data and then performs selection.

We first perform k -means clustering [28] on the target unlabeled dataset to obtain an initial clustering, and then compute a per-cluster mean feature vector:

$$\bar{\mathbf{h}}_c^u = \frac{1}{N_c^u} \sum_{\substack{(\mathbf{x}, \hat{y}) \in \mathcal{D}^u \\ \hat{y} = c}} f(\mathbf{x}), \quad (10)$$

where \hat{y} is the clustering label generated by running k -means on the unlabeled target dataset \mathcal{D}^u , and N_c^u is the number of images in cluster c . As in Eq. (7), we can directly use the labels to generate the per-category vectors $\bar{\mathbf{h}}_c^l$ for \mathcal{D}^l .

The key for the data selection is to generate lower weights for data that are too dissimilar or too similar. Our binning approach achieves this by first discarding dissimilar samples and then removing data that is too similar. To discard dissimilar samples, we set a similarity threshold to filter out distant unrelated samples. The threshold is determined by computing the similarities between the data points *within* the target unlabeled data. This gives us a measure of how similar the categories in the target unlabeled data should be, and thus we can filter out data that is less similar to the target data than this threshold. Next, we need to remove data that is too *similar*. To do this, we rank the labeled categories in the remaining source data based on their similarity to the target data and then select the labeled categories that are the most dissimilar from the target unlabeled data. This operation removes data that is too similar to the target data. The above selection is performed by setting the weight factor to a ‘hard’ value, *i.e.* $\omega \in \{0, 1\}$. In practice, we can discard the samples with weight $\omega_i = 0$ during training as they do not directly impact the loss. A more detailed description is provided in the supplementary material.

Soft weighting. The issue with the previous approach is that the choice of ‘hard’ binning enforces a hard selection, whereby some labeled data is kept and the rest is discarded. Additionally, the binning method requires performing clustering on the target unlabelled data first, which requires knowledge (or an estimate) of how many clusters there should be in the unlabelled target data. An alternative approach is to use soft weights to downweight data that is either too similar or too dissimilar. This intuition motivates us to introduce a new approach for soft selection that leverages the functional form of the beta distribution. Specifically, the probability density function of the beta distribution is defined as:

$$b(x; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1}, \quad (11)$$

where $B(\alpha, \beta)$ is the beta function which ensures the distribution sums to 1, $\alpha, \beta > 0$ are shape parameters, and $0 \leq x \leq 1$ is the input similarity.

We can use this distribution as a mechanism to generate weights. Specifically, we can input a similarity score in the range $[-1, 1]$ to $b(\cdot; \alpha, \beta)$, and use the output as the weight for the respective data point from the labeled source. Using different shape parameters with for the distribution changes which data is weighted higher or lower. We compute the similarity score as the cosine similarity of each of labeled centroids to the farthest unlabeled data point:

$$\gamma_c^l = \min_{\mathbf{x} \in \mathcal{D}^u} \cos(\bar{\mathbf{h}}_c^l, f(\mathbf{x})), \quad (12)$$

where $\cos(\cdot, \cdot)$ denotes the cosine similarity function. The intuition for choosing the farthest unlabeled data point is to see how related the labeled category is to the farthest unlabeled data, an ablation is provided in the supplementary for this design choice. Compared to the selection used in [9] and our hard binning method, our soft weighting approach is more lightweight as it does not require running a clustering first on the unlabeled data. The weight of the labeled data is thus calculated as:

$$\omega^c = b(\gamma_c^l; \alpha, \beta), \quad (13)$$

where ω^c is a weight value that is assigned to all data points from the category c , *i.e.* $\omega_i = \omega^c, \forall (\mathbf{x}_i, y) \in \mathcal{D}^l$ and $y = c$. This formulation gives us a simple way of generating smaller weights for data points that are too similar or dissimilar. Additionally, this formulation allows us to further tune the shape of the curve by changing the shape parameters α and β , as shown in Fig. 2 left. For example, setting both α and β to 1 yields a degenerated solution of uniform weight (as in conventional discovery). When $\alpha, \beta > 1$ and $\alpha > \beta$ the algorithm will favor more similar data points to the target (as in [9]), when $\alpha, \beta > 1$ and $\beta > \alpha$ it will favor more dissimilar ones, and when $\alpha, \beta > 1$ and $\beta = \alpha$ it will favor data points that are not too similar or not too dissimilar.

4 Experiments

Here, we quantify the impact of labeled data on category discovery performance and present the experimental results of our proposed data selection approaches.

4.1 Implementation details

For the majority of our experiments, we use the recent state-of-the-art generalized category discovery method SimGCD [49]. We use the DINOv2 [32] pre-trained ViT-Base/14 model [11] for our backbone $f(\cdot)$. Following the standard practice in category discovery, we freeze the backbone model up to the last block to avoid overfitting [45, 49]. As a baseline, we also report the discovery performance resulting from running ‘ k -means’ clustering on the DINOv2 features using the known number of categories for k . We evaluate on the commonly used Semantic-Shifts Benchmarks (SSB) [46]. SSB contains data from CUB [48], Stanford-Cars [25], and FGVC-Aircraft [29], and provides the category splits of the data for category discovery with an annotation of ‘Easy’, ‘Medium’, and ‘Hard’ categories that are built on the hierarchy from the dataset labels. The ‘Easy’ categories are those that can be easily separated from the labeled categories, similar for ‘Medium’ and ‘Hard’ categories. In this paper, we denote ‘Hard’ categories as ‘Similar’ and ‘Easy’ categories as ‘Dissimilar’. We use these annotations in the evaluation and motivation of the data selection problem. Additionally, we also leverage the Stanford-Dogs [23] dataset and a subset of Insect categories from iNat2021 [43]. The Stanford-Dogs dataset is also split into ‘Similar’, ‘Medium’, and ‘Dissimilar’ categories using its label hierarchy.

Evaluation protocol. To evaluate model performance, we adopt the commonly used clustering accuracy metric [45]. During the evaluation, given the ground truth labels y and the model predicted labels \tilde{y} , we calculate the accuracy as

$\text{ACC} = \frac{1}{M} \sum_{i=1}^M \mathbf{1}(y_i^* = \mathcal{P}(\tilde{y}_i))$, where $M = |\mathcal{D}^u|$, and \mathcal{P} is the optimal permutation that matches the predicted cluster assignments to the ground truth class labels. Additionally, we report ‘All’, ‘New’, and ‘Old’ performance as is common in the literature. ‘All’ is for the performance on all clusters, and ‘New’ and ‘Old’ denote the performance on novel and seen categories respectively. We report the average of three different runs in each table.

4.2 Evaluation

Table 1: Performance on ‘New’ categories when using different labeled data subsets to supervise SimGCD [49].

Labeled data	CUB	SCars	Aircraft	SDogs
k -means	32.1	13.8	16.8	30.2
Orig. data	57.7	45.0	51.8	58.2
OOD iNat-Insect	3.4	3.1	4.5	2.1
Similar	47.6	43.4	40.0	45.6
Medium	79.7	<u>43.2</u>	56.6	53.1
Dissimilar	<u>61.4</u>	38.4	<u>47.6</u>	<u>50.2</u>

‘Similar’, ‘Medium’, and ‘Dissimilar’ splits are based on the ground truth hierarchical labels of the dataset, thus the results are only for demonstration purposes as in the real world, we cannot obtain this ground truth information. We observe that the choice of labeled data has a large impact on discovery performance.

Perhaps surprisingly, the source split that performs the best (‘Medium’) is a subset of the source data that contains categories that are not too similar nor too dissimilar. Using the ‘Medium’ labeled data alone can match, or even outperform, using all labeled data in some cases (*i.e.* on CUB and Aircraft), despite using only one-third of the data. By taking a closer look at the source of errors in Fig. 3, we identify that when the model is using the ‘Similar’ splits, the majority of errors stem from misclassifying the novel categories to previously seen ones. When the model is trained with ‘Dissimilar’ splits, the error is mainly due to misclassifying the novel categories to incorrect novel clusters. When the labeled split comes from Out-of-Domain (OOD) insects, the error is almost entirely from misclassifying novel categories as incorrect novel clusters. These results indicate that when the model is trained with ‘Dissimilar’ or OOD data, it can learn to predict the novel categories as new clusters, but cannot learn

In this section, we demonstrate our motivating experiments and report the impact of labeled data selection on category discovery performance. Unless noted otherwise, the scores are the performance on ‘New’ categories.

Labeled data selection. In Tab. 1, we report the discovery performance of SimGCD [49] on each of the target datasets when using different source labeled data. Note that these ‘Similar




Fig. 3: Percentage of errors on three datasets when using ‘Similar’, ‘Medium’, ‘Dissimilar’, and ‘OOD’ labeled training sets. ‘Misclassified as New’ means that an example from a ‘New’ category is assigned to another ‘New’ one, and ‘Misclassified as Old’ means it is assigned to an ‘Old’ one.

Table 2: Category discovery results on the Semantic Shift Benchmark [46]. Here, labeled data is selected from **all** the datasets combined, *i.e.* we combined all the labeled data from the four datasets and the OOD insects. Performance change with respect to no-selection in presented in parentheses.

Methods	CUB			Stanford Cars			FGVC-Aircraft			Stanford Dogs		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New
<i>k</i> -means [28]	34.7	49.8	34.2	16.7	14.5	17.1	17.9	16.5	19.2	21.5	22.6	20.1
GCD [45]	44.8	50.9	41.2	35.9	55.2	26.7	41.3	38.7	45.8	48.7	54.1	46.1
GCD w/ [9]	40.7 (-4.1)	46.8 (-4.1)	36.8 (-4.4)	35.0 (-0.9)	55.9 (+0.7)	25.7 (-1.0)	34.8 (-6.5)	30.1 (-8.6)	36.9 (-8.9)	44.7 (-4.0)	50.6 (-3.5)	42.7 (-3.4)
GCD w/ Bins	47.9 (+3.1)	52.8 (+1.9)	44.2 (+3.0)	36.0 (+0.1)	55.0 (-0.2)	25.0 (-1.7)	44.9 (+3.6)	40.8 (+2.1)	47.2 (+1.4)	49.4 (+0.7)	54.0 (-0.1)	48.2 (+2.1)
GCD w/ Beta	47.6 (+2.8)	53.0 (+2.1)	43.7 (+2.5)	34.9 (-1.0)	56.2 (+1.0)	25.5 (-1.2)	44.1 (+2.8)	39.7 (+1.0)	46.2 (+0.4)	49.9 (+1.2)	55.5 (+1.4)	47.4 (+1.3)
XCon [14]	46.7	50.4	44.3	38.7	55.9	28.7	43.1	40.7	47.9	50.0	54.4	48.0
XCon w/ [9]	42.8 (-3.9)	45.7 (-4.7)	39.6 (-4.7)	39.2 (+0.5)	55.0 (-0.9)	26.9 (-1.8)	35.8 (-7.3)	31.0 (-9.7)	40.8 (-7.1)	45.9 (-4.1)	52.0 (-2.4)	42.2 (-5.8)
XCon w/ Bins	48.2 (+1.5)	51.0 (+0.6)	46.7 (+2.4)	39.8 (+1.1)	55.7 (-0.2)	27.9 (-0.8)	44.8 (+1.7)	42.1 (+1.4)	48.3 (+0.4)	52.3 (+2.3)	56.1 (+1.7)	48.9 (+0.9)
XCon w/ Beta	47.8 (+1.1)	51.8 (+1.4)	46.0 (+1.7)	39.5 (+0.8)	56.6 (+0.7)	28.9 (+0.2)	45.0 (+1.9)	42.6 (+1.9)	49.3 (+1.4)	52.5 (+2.5)	56.3 (+1.9)	49.4 (+1.4)
SimGCD [49]	55.1	61.2	52.3	50.4	68.3	40.9	51.8	56.2	49.3	57.2	60.1	54.2
SimGCD w/ [9]	50.3 (-4.8)	58.3 (-2.9)	46.5 (-5.8)	51.6 (+1.2)	69.4 (+1.1)	44.7 (+3.8)	44.2 (-7.6)	48.7 (-7.5)	40.7 (-8.6)	55.2 (-2.0)	57.2 (-2.9)	50.1 (-4.1)
SimGCD w/ Bins	58.2 (+3.1)	64.6 (+3.4)	56.7 (+4.4)	52.4 (+2.0)	68.2 (-0.1)	45.9 (+5.0)	55.7 (+3.9)	57.9 (+1.2)	56.8 (+7.5)	58.1 (+0.9)	62.3 (+2.2)	55.6 (+1.4)
SimGCD w/ Beta	58.6 (+3.5)	64.5 (+3.3)	56.8 (+4.5)	52.8 (+2.4)	68.9 (+0.6)	46.3 (+5.4)	56.3 (+4.5)	58.3 (+2.1)	56.9 (+7.6)	58.4 (+1.2)	62.8 (+2.7)	56.6 (+2.4)
μ GCD [47]	59.3	67.8	58.4	58.6	72.3	48.3	55.9	63.5	54.2	63.4	65.1	58.3
μ GCD w/ [9]	54.3 (-5.0)	62.4 (-5.4)	53.5 (-4.9)	54.6 (-4.0)	68.1 (-4.2)	45.1 (-3.2)	49.1 (-6.8)	56.9 (-6.6)	48.7 (-7.5)	56.2 (-7.2)	59.8 (-5.3)	54.1 (-4.2)
μ GCD w/ Bins	61.2 (+1.9)	69.3 (+1.5)	60.1 (+1.7)	59.5 (+0.9)	73.5 (+1.2)	50.4 (+2.1)	57.2 (+1.3)	64.5 (+1.0)	55.4 (+1.2)	64.5 (+1.1)	66.3 (+1.2)	59.0 (+1.3)
μ GCD w/ Beta	62.3 (+3.0)	70.5 (+2.7)	61.8 (+3.4)	60.4 (+1.8)	74.3 (+2.0)	51.0 (+2.7)	58.2 (+2.3)	65.8 (+2.3)	57.7 (+3.5)	66.4 (+3.0)	67.8 (+2.7)	61.3 (+3.0)

Table 3: Discovery results on NABirds [42], Herbarium-19 [40], and ImageNet [35]. The source data consists of these three datasets plus the source data used in Tab. 4.

Method	NABirds			Herb-19			ImageNet-100			ImageNet-1k-SSB		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New
SimGCD	54.6	61.2	49.0	44.2	57.6	37.0	78.5	86.7	73.5	50.2	73.6	42.6
SimGCD w/ [9]	50.1 (-4.5)	54.6 (-6.6)	44.8 (-4.2)	38.1 (-4.1)	52.6 (-5.0)	38.7 (+1.7)	76.4 (-2.1)	84.2 (-2.5)	70.2 (-3.3)	48.2 (-2.0)	68.9 (-4.7)	40.2 (-2.4)
SimGCD w/ Bins	56.2 (+1.4)	63.5 (+2.3)	51.6 (+2.6)	44.7 (+0.5)	58.1 (+0.5)	36.8 (-0.2)	78.2 (-0.3)	87.3 (+0.6)	73.3 (-0.2)	51.4 (+1.2)	73.5 (-0.1)	43.7 (+1.1)
SimGCD w/ Beta	55.7 (+1.1)	61.8 (+1.6)	50.8 (+1.8)	45.2 (+1.0)	57.8 (+0.2)	38.3 (+1.3)	78.9 (+0.4)	87.0 (+0.3)	73.4 (-0.1)	51.7 (+1.5)	74.0 (+0.4)	44.0 (+1.4)

a discriminative enough representation to separate the novel categories apart. On the other hand, when the model is trained with ‘Similar’ data, it struggles to separate novel categories from seen ones, resulting in a higher error when misclassifying novel categories as seen.

Selecting from a large pool of data. Previous category discovery works explore the discovery problem on a target unlabeled dataset using a fixed labeled dataset containing the same visual categories. Instead, we consider a more realistic scenario where we need select a suitable labeled dataset from a large repository of potentially unrelated datasets. Only some of the categories in this repository will be relevant to the unlabeled target. In our experiments, we set the pool of datasets to be the combination of the labeled data from CUB, Stanford-Cars, Stanford-Dogs, FGVC-Aircraft, and iNat-Insect, and then evaluate discovery performance on a given target dataset. The results are presented in Tab. 2.

We observe that using the combined labeled data to train discovery models it actually performs worse than just using the original labeled data for a given target task (see Tab. 4), even if the combined labeled data is three-to-four times larger. This suggests that for category discovery, more labeled data may not always be helpful. We conduct experiments comparing the data selection method of [9] and our proposed approaches (‘w/ Bins’ and ‘w/ Beta’). We observe that [9], which selects similar data to the target, is not suitable for category discovery and it results in a performance drop in most cases. In contrast, discovery models trained with our data selection methods can match, or even outperform, models trained with the provided labeled data split, indicating their

Table 4: Category discovery results on the Semantic Shift Benchmark [46]. Here, labeled data for a given dataset is selected from the **original** labeled splits for that dataset, *e.g.* we only use labeled data from CUB for the experiments on CUB.

Methods	CUB			Stanford Cars			FGVC-Aircraft			Stanford Dogs		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New
<i>k</i> -means [28]	34.3	38.9	32.1	12.8	10.6	13.8	16.0	14.4	16.8	18.9	19.1	18.0
GCD [45]	51.3	56.6	48.7	39.0	57.6	29.9	45.0	41.1	46.9	54.6	58.7	50.1
GCD w/ [9]	45.7 (-5.6)	51.2 (-5.4)	44.7 (-4.0)	37.2 (-1.8)	56.0 (-1.6)	23.7 (-6.2)	40.1 (-4.9)	38.7 (-2.4)	40.0 (-6.9)	48.9 (-5.7)	50.2 (-8.5)	45.7 (-4.4)
GCD w/ Bins	50.2 (-1.1)	54.7 (-1.9)	49.7 (+3.0)	40.3 (+1.3)	57.7 (+0.1)	32.0 (+2.1)	39.7 (-5.3)	39.9 (-1.2)	41.3 (-5.6)	53.8 (-0.8)	54.3 (-4.4)	50.1 (+0.0)
GCD w/ Beta	52.3 (+1.0)	57.8 (+1.2)	49.5 (+0.8)	41.2 (+2.2)	59.0 (+1.4)	33.4 (+3.5)	44.7 (-0.3)	41.3 (+0.2)	46.4 (-0.5)	55.4 (+0.8)	59.5 (+0.8)	50.6 (+0.5)
XCon [14]	52.1	54.3	51.0	40.5	58.8	31.7	47.7	44.4	49.4	56.1	59.1	52.4
XCon w/ [9]	48.3 (-3.8)	50.2 (-4.1)	47.8 (-3.2)	39.8 (-0.7)	57.2 (-1.6)	30.2 (-1.5)	43.5 (-4.2)	40.3 (-4.1)	45.2 (-4.2)	54.1 (-2.0)	56.7 (-2.4)	49.7 (-2.7)
XCon w/ Bins	51.7 (-0.4)	52.8 (-1.5)	53.7 (+2.7)	40.8 (+0.3)	57.1 (-1.7)	34.9 (+3.2)	48.3 (+0.6)	44.0 (-0.4)	49.7 (+0.3)	58.9 (+2.8)	61.2 (+2.1)	54.0 (+2.2)
XCon w/ Beta	52.8 (+0.7)	55.5 (+1.2)	51.4 (+0.4)	41.7 (+1.2)	59.9 (+1.1)	35.9 (+4.2)	49.4 (+1.7)	46.1 (+1.7)	51.3 (+1.9)	60.0 (+3.9)	62.3 (+3.2)	55.0 (+2.0)
SimGCD [49]	60.3	65.6	57.7	53.8	71.9	45.0	54.2	59.1	51.8	62.8	64.7	58.2
SimGCD w/ [9]	58.9 (-1.4)	63.6 (-2.0)	55.4 (-2.3)	54.9 (+1.1)	72.8 (+0.9)	46.7 (+1.7)	48.9 (-5.3)	56.3 (-1.8)	44.3 (-7.5)	60.2 (-2.6)	65.1 (+0.4)	50.9 (-7.3)
SimGCD w/ Bins	63.2 (+3.9)	66.7 (+1.1)	62.8 (+5.1)	53.0 (-0.8)	68.7 (-3.2)	44.5 (0.5)	59.8 (+5.6)	63.5 (+4.4)	56.0 (+4.2)	64.7 (+1.9)	66.0 (+1.3)	58.8 (+0.6)
SimGCD w/ Beta	64.5 (+5.2)	67.9 (+2.2)	64.0 (+6.3)	54.1 (+0.3)	70.1 (-1.8)	46.3 (+1.8)	60.7 (+6.5)	64.8 (+5.7)	57.1 (+3.3)	66.0 (+3.2)	67.1 (+2.4)	59.0 (+1.4)
μ GCD [47]	65.7	68.0	64.6	58.7	76.5	50.2	58.7	64.7	56.8	67.9	69.4	62.5
μ GCD w/ [9]	64.3 (-1.4)	65.2 (-2.8)	60.4 (-4.2)	57.1 (-1.7)	75.0 (-1.5)	48.7 (-1.3)	54.9 (-3.8)	63.5 (-1.2)	52.3 (-4.5)	64.2 (-3.7)	66.1 (-3.3)	58.4 (-4.1)
μ GCD w/ Bins	67.3 (+1.6)	69.0 (+1.0)	65.7 (+1.1)	60.4 (+1.3)	77.8 (+1.3)	51.5 (+1.3)	60.3 (+1.6)	65.9 (+1.1)	58.9 (+2.1)	69.3 (+1.4)	71.1 (+1.7)	64.2 (+1.7)
μ GCD w/ Beta	68.4 (+2.7)	70.1 (+2.1)	66.6 (+2.0)	61.0 (+1.9)	78.6 (+2.1)	52.2 (+2.0)	62.0 (+2.3)	67.2 (+2.4)	59.8 (+3.0)	70.3 (+2.4)	72.0 (+2.6)	65.4 (+2.9)

effectiveness. Furthermore, in Tab. 3, we present results on larger and more challenging datasets. We observe that our methods still yields better performance compared to the vanilla training baseline, and the improvements on these challenging benchmarks are significant above noise. Interestingly, on the NABirds dataset the binning method outperforms the beta soft weighting. We conjecture this to the very fine-grained nature of this dataset, thus all labeled data has roughly same distance to the target unlabeled data, resulting in similar weights for all labeled data.

Selecting from the original labeled data. In Tab. 4, we present the results of data selection using the originally provided labeled data for each dataset. This is the standard setting explored in existing discovery works. The most interesting observation here comes from comparing our data selection methods with the standard approach of directly using all data of the provided labeled dataset. We can see that in almost all cases, using our soft weighting selection method (‘w/ Beta’) outperforms this baseline. Owing to our improvement over the already highly performant μ GCD [47], our results set a new state-of-the-art on each dataset. This result further strengthens our claim that labeled data matters for category discovery, as even the original provided labeled data is not the optimal one for category discovery. Thus, more effort should be put into developing new methods for automatically selecting suitable labeled data subsets for target datasets of interest. We also observe that [9] is not effective at data selection in the category discovery setting. In almost all cases, it decreases performance as it is designed to select the most *similar* data to the target.

4.3 Ablations

In this section, we present ablation experiments on our proposed method with the beta distribution. More ablations can be found in the supplementary.

Different of α, β values. In Sec. 3.3, we showed that by changing the value of α and β , we can control the shape of the curve for generating data weights thus interpolating between different strategies. In Tab. 5 we observe that when

setting the values to $\alpha = 5$ and $\beta = 1$, to favor more similar samples like in [9], the performance drops. Setting both α and β values to larger than 1 yields consistent performance improvements, and our default choice of $\alpha = \beta = 5$ yields consistently strong performance across the three datasets.

Hard threshold cutting. In Sec. 3.3, we discussed using the weight generated by Eq. (13) to weigh the influence of each example. Tab. 6 shows the ablation of thresholding the generated weight to $\{0, 1\}$. By thresholding the weight ω_i of an instance \mathbf{x}_i to 0, we can effectively remove that sample from the training dataset. Thus the training time can be reduced. These results demonstrate that while using the soft weight performs the best, removing some examples from training not only does not hurt performance but improves it to some extent.

Table 5: Ablations on different values of α and β . We default to $\alpha = \beta = 5$ for our experiments. Scores are performance on ‘New’ categories. **Table 6:** Our main experiments use soft weighting, here we ablate hard thresholding those weights. Scores are performance on ‘New’ categories.

α β	CUB	Aircraft	SDogs
1 1	52.3	49.3	54.2
5 1	44.5 (-7.8)	39.5 (-9.8)	48.9 (-5.3)
5 3	57.5 (+5.2)	56.4 (+7.1)	56.2 (+2.0)
3 5	57.3 (+5.0)	56.7 (+7.4)	56.7 (+2.5)
5 5	56.8 (+4.5)	56.9 (+7.6)	56.6 (+2.4)

	CUB	Aircraft	SDogs
baseline	52.3	49.3	54.2
threshold=0.2	55.0 (+2.7)	55.8 (+6.5)	56.0 (+1.8)
threshold=0.5	54.7 (+2.4)	54.6 (+4.9)	55.4 (+1.2)
threshold=1.0	53.4 (+1.1)	54.0 (+4.3)	55.1 (+0.9)
soft-weight	56.8 (+4.5)	56.9 (+7.6)	56.6 (+2.4)

4.4 Limitations

One limitation of our proposed data selection method is that it requires specifying some hyperparameters (*e.g.* the parameters of the beta distribution). Notably, we use the same hyperparameters across all datasets and demonstrate the robustness of our method to different choices. There is also an additional cost incurred by our approach associated with computing the source data training weights. However, this is negligible compared to the overall cost of training. Like all discovery methods, if the resulting discovered categories were not verified by a human expert, there could be potential negative impact of our work if it was applied in safety-critical applications.

5 Conclusion

We explored the problem of category discovery through the lens of labeled data selection. Perhaps surprisingly, we showed that the precise choice of labeled data matters significantly for discovery performance. We demonstrated that the most effective labeled source is not too dissimilar to the unlabeled data as to provide no useful signal, nor is it too similar as to cause confusion between novel and previously seen similar categories. We presented a new approach for automatically weighting different instances from the labeled dataset which results in improved discovery performance compared to standard splits from the existing literature. Interestingly, we showed that this improvement is not the result of simply selecting more labeled data, but in fact less, *i.e.* less is more. We hope that our work opens up a new line of enquiry into data selection for category discovery.

Acknowledgements. This work was in part supported by a Royal Society Research Grant, the Edinburgh–Copenhagen Strategic Partnership Seed-Fund, and the Pioneer Centre for AI, DNRF grant number P1.

References

1. Arthur, D., Vassilvitskii, S., et al.: k-means++: The advantages of careful seeding. In: Symposium on Discrete Algorithms (2007) [24](#)
2. Azizpour, H., Razavian, A.S., Sullivan, J., Maki, A., Carlsson, S.: Factors of transferability for a generic convnet representation. PAMI (2015) [4](#)
3. Bao, Y., Li, Y., Huang, S.L., Zhang, L., Zheng, L., Zamir, A., Guibas, L.: An information-theoretic approach to transferability in task transfer learning. In: ICIP (2019) [4](#)
4. Bolya, D., Mittapalli, R., Hoffman, J.: Scalable diverse model selection for accessible transfer learning. In: NeurIPS (2021) [4](#)
5. Cao, K., Brbić, M., Leskovec, J.: Open-world semi-supervised learning. In: ICLR (2022) [3](#), [5](#)
6. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: ECCV (2018) [1](#)
7. Chi, H., Liu, F., Yang, W., Lan, L., Liu, T., Han, B., Niu, G., Zhou, M., Sugiyama, M.: Meta discovery: Learning to discover novel classes given very limited data. In: ICLR (2022) [3](#)
8. Cole, E., Wilber, K., Van Horn, G., Yang, X., Forni, M., Perona, P., Belongie, S., Howard, A., Mac Aodha, O.: On label granularity and object localization. In: ECCV (2022) [4](#)
9. Cui, Y., Song, Y., Sun, C., Howard, A., Belongie, S.: Large scale fine-grained categorization and domain-specific transfer learning. In: CVPR (2018) [2](#), [4](#), [7](#), [8](#), [10](#), [12](#), [13](#), [14](#)
10. Dehghani, M., Djolonga, J., Mustafa, B., Padlewski, P., Heek, J., Gilmer, J., Steiner, A.P., Caron, M., Geirhos, R., Alabdulmohsin, I., et al.: Scaling vision transformers to 22 billion parameters. In: ICML (2023) [4](#)
11. 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) [10](#)
12. Du, R., Chang, D., Liang, K., Hospedales, T., Song, Y.Z., Ma, Z.: On-the-fly category discovery. In: CVPR (2023) [3](#)
13. Fang, A., Jose, A.M., Jain, A., Schmidt, L., Toshev, A., Shankar, V.: Data filtering networks. In: ICLR (2024) [4](#)
14. Fei, Y., Zhao, Z., Yang, S., Zhao, B.: Xcon: Learning with experts for fine-grained category discovery. In: BMVC (2022) [3](#), [5](#), [12](#), [13](#), [21](#), [23](#), [24](#)
15. Fini, E., Sangineto, E., Lathuilière, S., Zhong, Z., Nabi, M., Ricci, E.: A unified objective for novel class discovery. In: ICCV (2021) [3](#)
16. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: CVPR (2018) [25](#)
17. Guo, S., Huang, W., Zhang, X., Srikhanta, P., Cui, Y., Li, Y., Adam, H., Scott, M.R., Belongie, S.: The imaterialist fashion attribute dataset. In: ICCV Workshops (2019) [4](#)
18. Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Automatically discovering and learning new visual categories with ranking statistics. In: ICLR (2020) [3](#)

19. Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Autonovel: Automatically discovering and learning novel visual categories. *PAMI* (2021) [3](#)
20. Han, K., Vedaldi, A., Zisserman, A.: Learning to discover novel visual categories via deep transfer clustering. In: *ICCV* (2019) [1](#), [3](#)
21. Ji, X., Henriques, J.F., Vedaldi, A.: Invariant information clustering for unsupervised image classification and segmentation. In: *ICCV* (2019) [1](#)
22. Katharopoulos, A., Fleuret, F.: Not all samples are created equal: Deep learning with importance sampling. In: *ICML* (2018) [4](#)
23. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: *FGVC Workshop at CVPR* (2011) [10](#), [23](#)
24. Kim, H., Suh, S., Kim, D., Jeong, D., Cho, H., Kim, J.: Proxy anchor-based unsupervised learning for continuous generalized category discovery. In: *ICCV* (2023) [3](#)
25. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: *ICCV Workshops* (2013) [10](#), [23](#)
26. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *ECCV* (2014) [4](#)
27. Loshchilov, I., Hutter, F.: Online batch selection for faster training of neural networks. In: *ICLR Workshops* (2016) [4](#)
28. MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* (1967) [8](#), [12](#), [13](#)
29. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. *arXiv:1306.5151* (2013) [10](#), [20](#), [23](#)
30. Mindermann, S., Brauner, J.M., Razzak, M.T., Sharma, M., Kirsch, A., Xu, W., Hölting, B., Gomez, A.N., Morisot, A., Farquhar, S., et al.: Prioritized training on points that are learnable, worth learning, and not yet learnt. In: *ICML* (2022) [4](#)
31. Nguyen, C., Hassner, T., Seeger, M., Archambeau, C.: Leap: A new measure to evaluate transferability of learned representations. In: *ICML* (2020) [4](#)
32. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *TMLR* (2024) [4](#), [10](#)
33. Pu, N., Zhong, Z., Sebe, N.: Dynamic conceptional contrastive learning for generalized category discovery. In: *CVPR* (2023) [3](#), [5](#)
34. Rastegar, S., Doughty, H., Snoek, C.: Learn to categorize or categorize to learn? self-coding for generalized category discovery. In: *NeurIPS* (2023) [3](#)
35. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *IJCV* (2015) [4](#), [12](#), [23](#)
36. Shao, D., Zhao, Y., Dai, B., Lin, D.: Finegym: A hierarchical video dataset for fine-grained action understanding. In: *CVPR* (2020) [4](#)
37. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: *ICCV* (2017) [4](#)
38. Sun, Y., Shi, Z., Li, Y.: A graph-theoretic framework for understanding open-world semi-supervised learning. In: *NeurIPS* (2023) [3](#)
39. Sun, Y., Shi, Z., Liang, Y., Li, Y.: When and how does known class help discover unknown ones? provable understanding through spectral analysis. In: *ICML* (2023) [3](#)
40. Tan, K.C., Liu, Y., Ambrose, B., Tulig, M., Belongie, S.: The herbarium challenge 2019 dataset. In: *FGVC Workshop at CVPR* (2019) [12](#), [23](#)

41. Tran, A.T., Nguyen, C.V., Hassner, T.: Transferability and hardness of supervised classification tasks. In: ICCV (2019) [4](#)
42. Van Horn, G., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeiritis, P., Perona, P., Belongie, S.: Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In: CVPR (2015) [4](#), [12](#), [23](#)
43. Van Horn, G., Cole, E., Beery, S., Wilber, K., Belongie, S., Mac Aodha, O.: Benchmarking representation learning for natural world collections. In: CVPR (2021) [10](#)
44. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: CVPR (2018) [4](#), [23](#)
45. Vaze, S., Han, K., Vedaldi, A., Zisserman, A.: Generalized category discovery. In: CVPR (2022) [1](#), [3](#), [5](#), [10](#), [12](#), [13](#), [21](#), [23](#), [24](#)
46. Vaze, S., Han, K., Vedaldi, A., Zisserman, A.: Open-set recognition: A good closed-set classifier is all you need? In: ICLR (2022) [10](#), [12](#), [13](#), [20](#)
47. Vaze, S., Vedaldi, A., Zisserman, A.: No representation rules them all in category discovery. In: NeurIPS (2023) [3](#), [5](#), [12](#), [13](#), [23](#), [25](#)
48. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: Caltech-UCSD Birds 200. Computation & Neural Systems Technical Report (2010) [10](#), [20](#), [23](#)
49. Wen, X., Zhao, B., Qi, X.: Parametric classification for generalized category discovery: A baseline study. In: ICCV (2023) [3](#), [5](#), [10](#), [11](#), [12](#), [13](#), [21](#), [25](#)
50. Xu, H., Xie, S., Tan, X.E., Huang, P.Y., Howes, R., Sharma, V., Li, S.W., Ghosh, G., Zettlemoyer, L., Feichtenhofer, C.: Demystifying clip data. In: ICLR (2024) [4](#)
51. Zhang, S., Khan, S., Shen, Z., Naseer, M., Chen, G., Khan, F.: Promptcal: Contrastive affinity learning via auxiliary prompts for generalized novel category discovery. In: CVPR (2023) [3](#)
52. Zhang, X., Jiang, J., Feng, Y., Wu, Z.F., Zhao, X., Wan, H., Tang, M., Jin, R., Gao, Y.: Grow and merge: A unified framework for continuous categories discovery. In: NeurIPS (2022) [1](#), [3](#)
53. Zhao, B., Han, K.: Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. In: NeurIPS (2021) [3](#), [4](#)
54. Zhao, B., Mac Aodha, O.: Incremental generalized category discovery. In: ICCV (2023) [1](#), [3](#), [4](#)
55. Zhao, B., Wen, X., Han, K.: Learning semi-supervised gaussian mixture models for generalized category discovery. In: ICCV (2023) [5](#), [21](#)
56. Zhong, Z., Fini, E., Roy, S., Luo, Z., Ricci, E., Sebe, N.: Neighborhood contrastive learning for novel class discovery. In: CVPR (2021) [3](#)