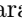






# SelEx: Self-Expertise in Fine-Grained Generalized Category Discovery - Appendix

Sarah Rastegar<sup>1</sup>, Mohammadreza Salehi<sup>1</sup>, Yuki M. Asano<sup>1\*</sup>,  
Hazel Doughty<sup>2</sup>, and Cees G. M. Snoek<sup>1</sup>

<sup>1</sup> University of Amsterdam

<sup>2</sup> Leiden University

This Appendix supplements the main text with additional details. Section 1 expands on related works, providing a comprehensive background. Section 2 contains proofs for the upper bounds mentioned in the main content. In Section 3, we detail our method’s pseudocode and elaborate on the components visually for better understanding. Section 4 includes further experimental results and ablation studies, demonstrating our method’s efficacy. The concluding discussion highlights future research avenues and acknowledges the limitations of our study, paving the way for further exploration.

## 1 Related Works

To provide a more in-depth analysis of the relevant literature, this appendix section delves into additional related works.

### 1.1 Self-Supervised Learning

Self-supervised learning has transformed the analysis of large-scale unlabeled data, learning rich representations. He *et al.* [26] introduced the concept of momentum contrast with MoCo, enhancing the quality of learned representations by utilizing two encoders and a contrastive loss to learn image features that distinguish between different views [26]. Following this, Chen *et al.* [12] simplified the self-supervised learning pipeline with SimCLR, by maximizing agreement between different augmented versions of the same image, using a contrastive loss function [12]. Caron *et al.* [7] introduced SwAV, which employs a clustering mechanism to enhance consistency between cluster assignments, thereby improving learning efficacy [7]. By employing a self-distillation with information noise injection, DINO [8] efficiently captures complex visual features without labeled data. DINOv2 [42] further refines this approach with key improvements, enhancing both feature quality and model adaptability.

### 1.2 Open-Set Recognition

The advent of open set recognition marked a significant milestone in the evolution of computational models, addressing the complexities of processing real-world data. This domain was first conceptualized by Scheirer *et al.* [55], who laid

---

\* Currently at University of Technology Nuremberg

the foundational theory for models capable of discerning between known and unknown data categories, a principle further developed by subsequent research [4, 54]. The pioneering application of deep learning techniques to tackle open-set recognition challenges was introduced through the development of OpenMax [5]. The core objective within open-set recognition is the accurate identification of known categories while effectively filtering out novel category samples. Various methodologies have been proposed to emulate the concept of “otherness” essential for distinguishing unknown categories. These approaches range from identifying significant reconstruction errors [46, 72], measuring the deviation from a set of predefined prototypes [10, 11, 56], to differentiating samples generated through adversarial processes [20, 31, 39, 74]. Despite its advances, a notable limitation of open-set recognition is its inclination to disregard all instances of novel classes, potentially omitting valuable information. This challenge underscores the ongoing quest for more sophisticated models capable of not only identifying the unknown but also accommodating the continuous expansion of the knowledge domain.

### 1.3 Novel Category Discovery

The foundation for this line of inquiry was laid by Han *et al.* [24], who adapted classification models to recognize novel categories based on knowledge from known categories. Initial strategies, as documented in several studies [22, 27, 28], typically employed a bifurcated approach. Initially, these methods focused exclusively on learning representations from annotated data, subsequently applying the acquired categorical structures to identify unknown categories in a distinct phase. However, the research trajectory has recently veered towards unified methodologies. For instance, studies like [18, 23, 36, 37, 52, 53, 73, 78, 81] exemplify this trend by integrating the learning process, simultaneously leveraging both labeled and unlabeled datasets to enhance representation learning. Despite the significant progress in this area, a critical issue remains unaddressed: the assumption that known and novel categories are mutually exclusive. This assumption undermines the performance of models in real-world settings, where instances may belong to either known or novel categories. This issue is what *Generalized Category Discovery* seeks to address.

### 1.4 Generalized Category Discovery

Generalized category discovery was introduced by Vaze *et al.* [59] and Cao *et al.* [6]. It provides models with unlabeled data from both novel and known categories, placing it within the realm of semi-supervised learning, a domain thoroughly investigated in the machine learning literature [9, 40, 44, 51, 70]. The unique challenge in generalized category discovery is handling categories without any labeled instances, which introduces additional complexity. There are primarily two approaches to address this challenge. The first employs a series of prototypes as reference points, *e.g.*, [2, 13, 14, 25, 29, 58, 63, 65, 68, 69, 75]. The second approach leverages local similarity as weak pseudo-labels for each sample and utilizes sample similarities to form local clusters [3, 13, 16, 19, 25, 43, 48, 49, 76, 79]. Studies

such as [60, 62, 65, 76] employ mean-teacher networks to tackle the issues arising from noisy pseudo-labels. Additionally, other research efforts leverage cues from alternative modalities to identify categories in a multimodal fashion [1, 45, 80]. Our approach introduces ‘self-expertise’, a novel concept aimed at hierarchical learning of known and unknown categories. This technique emphasizes the focus on the samples from identical clusters at each level. This method is particularly effective in overcoming the limited availability of positive samples per category, while also enhancing the identification of subtle differences among negative samples.


### 1.5 Hierarchical Representation Learning

In the realm of leveraging hierarchical categories for enhanced representation learning, several approaches have been introduced. Zhang *et al.* [77] employ multiple label levels to augment their models’ representational capacity through hierarchical contrastive learning. Similarly, Guo *et al.* [21] extract pseudo-labels to facilitate hierarchical contrastive learning, with a unique emphasis on ensuring signals remain positive within identical clusters. The hierarchical structure of categories has been explored in previous works such as [35, 67], aiming to enhance the identification of novel categories within the context of open-set recognition. Additionally, Rastegar *et al.* [50] utilize a hierarchical structure for multimodal data to infer missing modalities, which can be categories. An alternative perspective on extracting hierarchical structures of categories through self-supervision can be viewed through the framework of deep metric learning [30, 71]. Our research parallels these efforts by also utilizing hierarchical pseudo-labels. However, our methodology diverges by incorporating negative samples from the same cluster to foster generalized category discovery, setting our approach apart.

Moreover, Otholt *et al.* [43] and Banerjee *et al.* [3] have proposed hierarchical strategies aimed at generalized category discovery, primarily leveraging neighborhood structures to define refined categories with precision. Rastegar *et al.* [49] introduced a technique for learning an implicit category tree, which aids in the hierarchical self-coding of categories, ensuring category similarity is maintained throughout all levels of the hierarchy. Similarly, Wang *et al.* [64] leverage taxonomic context consistency to maintain coherence in the targeted labels, ensuring alignment across both coarse and fine-grained pseudo-labels. Our work distinguishes itself from these methodologies by employing weak supervision from samples within each hierarchy level to minimize the impact of misclassifications on lower levels. Furthermore, our approach emphasizes hard negatives within the same cluster for unsupervised self-expertise, thereby improving the model’s capability to identify subtle distinctions within fine-grained classifications.

## 2 Theory

In this section, we motivate our approach of hierarchical contrastive learning and why it is particularly well-suited for category discovery in unseen data.



**Fig. 1: Bayesian Network for the Generalized Category Discovery.**

### 2.1 Notations and Definitions

Let’s consider the simple Bayesian networks depicted in Fig. 1. Here,  $\mathbf{x}_i$  and  $\mathbf{x}_j$  are different samples or different views of the same sample which are observed. Their corresponding ground truth context variables  $\mathbf{c}_i$  and  $\mathbf{c}_j$ s are variables we aim to extract information about. These context variables can be partly observed in the form of labels, as is the case for supervised contrastive learning. However, these context variables are unobserved for unsupervised contrastive learning, which we show with  $z_i$  and  $z_j$ , respectively. Random variable  $y$  will indicate if its two parents have the same value, or in other terms; it will provide the contrastive labels.

Let’s assume that we have  $K$  total categories and the number of total samples as  $N$ . For each context variable  $c$ , we show the number of samples assigned to it by  $|c|$ . Hence,  $|z_i|$  indicates the number of samples that the model has assigned to modulo set  $\{z_i\}_K$  and  $|\mathbf{c}_i|$  indicates how many samples the true distribution has assigned to labels  $\{\mathbf{c}_i\}_K$ . Finally, we show the true distribution of samples with  $p$  and the approximated one with  $\hat{p}$ . In the next sections, we use labels instead of ground truth context variables, but note that these labels can have different hierarchy levels from the ground truth labels. To prevent ambiguity, we emphasize ground truth labels whenever we use them.

### 2.2 Problem Definition

For contrastive training, the goal is to estimate the true distribution of equality of the ground truth context variables. Hence, for samples  $i$  and  $j$ , the goal is to approximate the following true distribution,

$$p(y=1|\mathbf{c}_i, \mathbf{c}_j) = \mathbb{1}(\mathbf{c}_i=\mathbf{c}_j), \quad (1)$$

in which  $\mathbb{1}$  is the identity operator, which is one only when its inner condition holds and zero otherwise. Depending on the problem we aim to solve, in training, we have no access to the ground truth context variables as in unsupervised contrastive learning, or we have partial access as in supervised contrastive learning. In both contrastive learning formulations, we aim to minimize the KL divergence between the true distribution  $p$  and our model distribution  $\hat{p}$  which in case of unsupervised contrastive learning will be,

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{x}_i, \mathbf{x}_j)], \quad (2)$$

and for supervised contrastive learning will be,

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{c}_i, \mathbf{x}_j)]. \quad (3)$$

In the generalized category discovery problem, for labeled samples, we can consider the context variable as  $\mathbf{c}_i$  since it is observed, while for unlabeled samples, both context variables are unobserved.

### 2.3 Supervised Contrastive Learning

**Theorem 1.** *Consider two samples  $i$  and  $j$  from the Bayesian network Fig. 1. If we have a total of  $N$  samples and  $K$  categories and the dataset is balanced, we will have the following upper bound if only  $i$ 's label  $\mathbf{c}_i$  is known:*

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{c}_i, \mathbf{x}_j)] \leq \ln \frac{N}{K}. \quad (4)$$

*Proof.* According to the Bayesian network shown in Fig. 1, we have:

$$\hat{p}(y|\mathbf{c}_i, \mathbf{x}_j) = \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)\hat{p}(y|\mathbf{c}_i, z_j). \quad (5)$$

Note that since  $\sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)=1$  and  $p(y|\mathbf{c}_i, \mathbf{c}_j)$  is independent of  $z_j$ <sup>3</sup>, we can consider that

$$p(y|\mathbf{c}_i, \mathbf{c}_j) = \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)p(y|\mathbf{c}_i, \mathbf{c}_j). \quad (6)$$

Also, since KL divergence is convex, we can have the following inequality:

$$\begin{aligned} D_{\text{KL}}\left[\sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)\hat{p}(y|\mathbf{c}_i, z_j)\right] \\ \leq \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{c}_i, z_j)]. \quad (7) \end{aligned}$$

If we use equations Eqs. (5) to (7) on the left-hand side of the inequality Eq. (4), we will have the following:

$$\begin{aligned} D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{c}_i, \mathbf{x}_j)] &= D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)\hat{p}(y|\mathbf{c}_i, z_j)] \\ &= D_{\text{KL}}\left[\sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)\hat{p}(y|\mathbf{c}_i, z_j)\right] \\ &\leq \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j)D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{c}_i, z_j)], \end{aligned}$$

<sup>3</sup> Note that  $p$  is the true distribution and  $\hat{p}$  is the distribution that is derived from the Bayesian network.

For  $D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, z_j)]$ , two possible scenarios can happen, i.e.,  $\mathbf{c}_i = \mathbf{c}_j$  and  $\mathbf{c}_i \neq \mathbf{c}_j$ . Fig. 2 depicts these different scenarios in which the model makes a mistake in assigning the label. We call these two scenarios intra-class divergence and inter-class divergence, respectively. Note that since the dataset is balanced, we can consider that  $|z_j| = |\mathbf{c}_j| = \frac{N}{K}$ . We also assume  $\hat{p}(\mathbf{c}_j) \approx p(\mathbf{c}_j) = \hat{p}(z_j) = \frac{1}{K}$ .

**Intra-class Divergence.** If  $\mathbf{c}_i = \mathbf{c}_j$ , we will have  $p(y=0|\mathbf{c}_i, \mathbf{c}_j) = 0$ , according to Eq. (1). So we will have:

$$\begin{aligned} D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, \mathbf{c}_j, z_j=k)] = & \\ & - \ln \hat{p}(y=1|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \\ & \mathbb{1}(\hat{p}(y=1|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \neq 0). \end{aligned} \quad (8)$$

In this equation  $\mathbb{1}(\hat{p}(y=1|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \neq 0)$  ensures that only valid  $k$ s has been considered. We don't include this function in the next equations for simplicity but implicitly consider the condition it imposes on  $k$ . This scenario is depicted in Fig. 2a. Note that unlike [41], we consider  $\hat{p}(y)$  to be a Bernoulli distribution; this is because we only consider one sample at a time to see if this sample belongs to the same class or not. This is because, in their approach, categorical cross entropy forces the model to use one positive sample while considering all the remainder as negatives. But in the contrastive version, we can have multiple positives per each batch. Similar to [41], we can write this probability as follows:

$$\hat{p}(y=1|\mathbf{c}_i, \mathbf{c}_j, z_j=k) = \frac{\hat{p}(z_j=k|\mathbf{c}_i, \mathbf{c}_j)}{\hat{p}(z_j=k)} = \frac{\hat{p}(z_j=k|\mathbf{c}_i, \mathbf{c}_j)}{\sum_{l=1}^K \frac{\hat{p}(z_j=l|\mathbf{c}_i, \mathbf{c}_j)}{\hat{p}(z_j=l)}}. \quad (9)$$


Using Bayes theorem, we will have the following:

$$\hat{p}(z_j=k|\mathbf{c}_j) = \frac{|z_j=k|}{|\mathbf{c}_j|} \hat{p}(\mathbf{c}_j|z_j=k) = \hat{p}(\mathbf{c}_j|z_j=k). \quad (10)$$

We can rewrite Eq. (8) as:

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, \mathbf{c}_j, z_j)] = - \ln \overbrace{\frac{\hat{p}(z_j=k|\mathbf{c}_j)}{\hat{p}(z_j=k)}}^A + \ln \underbrace{\sum_{l=1}^K \frac{\hat{p}(z_j=l|\mathbf{c}_j)}{\hat{p}(z_j=l)}}_B. \quad (11)$$

To calculate  $A$ , since in Eq. (8), the condition is that  $\hat{p}(y=1|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \neq 0$ , we can conclude that  $\hat{p}(z_j=k|\mathbf{c}_j) > 0$ . At least one sample of  $\mathbf{c}_j$  or  $\mathbf{c}_i$  should have  $z_j=k$ ,



**Fig. 2: Different scenarios for two sample category variables in supervised contrastive learning.** The latent variable  $z_j$  is extracted from the model. Model does the red assignments, while blue assignments indicate the ground truth category variable. Since only  $\mathbf{c}_i$  is known, the model makes a mistake if: (a) Both  $i$  and  $j$  have the same labels and the model has assigned  $j$  to another label  $k$ . (b)  $i$  and  $j$  have different labels but  $j$  is assigned to the label  $k = \mathbf{c}_i$ .

or else this probability will be zero. So we have:  $\hat{p}(z_j=k|\mathbf{c}_j) \geq \frac{1}{|\mathbf{c}_j|} = \frac{K}{N}$ . Therefore, we can find an upper bound for  $A$ :

$$A \leq -\ln\left(\frac{K^2}{N}\right). \quad (12)$$

To calculate  $B$ ,

$$B = \ln\left(K \sum_{l=1}^K \hat{p}(z_j=l|\mathbf{c}_j)\right) = \ln K. \quad (13)$$

If we use equation Eqs. (7), (12) and (13) we have:

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, \mathbf{x}_j)] \leq \ln \frac{N}{K}. \quad (14)$$

**Inter-class Divergence** Here we have  $\mathbf{c}_j \neq \mathbf{c}_i$ , for this scenario which is depicted in Fig. 2b, we have a few changes from the previous part, particularly Eq. (8) will be as:

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, \mathbf{c}_j, z_j=k)] = -\ln \hat{p}(y=0|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \mathbb{1}(\hat{p}(y=0|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \neq 0). \quad (15)$$

Since we can consider  $\hat{p}(y=0|\mathbf{c}_i, \mathbf{c}_j, z_j=k) = 1 - \hat{p}(y=1|\mathbf{c}_i, \mathbf{c}_j, z_j=k)$ , Eq. (11) becomes

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, \mathbf{c}_j, z_j)] = \underbrace{-\ln \sum_{l \neq k} \frac{\hat{p}(z_j=l|\mathbf{c}_j)}{\hat{p}(z_j=l)}}_A + \underbrace{\ln \sum_{l=1}^K \frac{\hat{p}(z_j=l|\mathbf{c}_j)}{\hat{p}(z_j=l)}}_B. \quad (16)$$

To calculate  $A$

$$A = -\ln\left(K \sum_{l \neq k} \hat{p}(z_j=l|\mathbf{c}_j)\right) = -\ln K - \ln(1 - \hat{p}(z_j=k|\mathbf{c}_j)). \quad (17)$$

Again, since in Eq. (8), the condition is that  $\hat{p}(y=0|\mathbf{c}_i, \mathbf{c}_j, z_j=k) \neq 0$ , we can conclude that  $\hat{p}(z_j=k|\mathbf{c}_j) < 1$ . At least one sample of  $\mathbf{c}_j$  should have  $z_j=k$ , or else this probability will be zero. So we have:  $\hat{p}(z_j=k|\mathbf{c}_j) \leq (1 - \frac{1}{|\mathbf{c}_j|}) = 1 - \frac{K}{N}$ . Therefore, we can find a similar upper bound for  $A$ :

$$A \leq -\ln\left(\frac{K^2}{N}\right). \quad (18)$$

We have already calculated  $B$  in Eq. (13). If we use equation Eqs. (7), (13) and (17) we again have:


$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|\mathbf{c}_i, \mathbf{x}_j)] \leq \ln \frac{N}{K}. \quad (19)$$

Since in both cases, the upper bound holds the theorem is proven.  $\square$

## 2.4 Unsupervised Contrastive Learning

For the unsupervised scenario, both labels are unknown, this means that we only have access to the inputs  $\mathbf{x}_i$  and  $\mathbf{x}_j$ . This means that for unlabelled samples they will follow the Bayesian network shown in Fig. 3. We can state a similar theorem for the unsupervised case as follows:

**Theorem 2.** *Consider two samples  $i$  and  $j$  from the Bayesian network Fig. 1. If we have a total of  $N$  samples and  $K$  categories and the dataset is balanced, we will have the following upper bound if neither of  $i$  and  $j$  labels is known:*



**Fig. 3: Bayesian Network for the Generalized Category Discovery.** Shaded nodes are observed variables  $\mathbf{x}_i$ ,  $\mathbf{x}_j$  corresponds to images  $i$  and  $j$ , and  $\mathbf{c}_i$  and  $\mathbf{c}_j$  which are the ground-truth category variable.  $z_i$  and  $z_j$  are the latent category variables extracted from the model.

$$D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{x}_i, \mathbf{x}_j)] \leq \ln \frac{N}{K}. \quad (20)$$

*Proof.* According to the Bayesian network shown in Fig. 3, we have:

$$\hat{p}(y|\mathbf{x}_i, \mathbf{x}_j) = \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) \hat{p}(y|z_i, z_j). \quad (21)$$

Note that since  $\sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) = 1$  and  $\sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) = 1$  and  $p(y|\mathbf{c}_i, \mathbf{c}_j)$  is independent of  $z_i$  and  $z_j$ , we can consider that

$$p(y|\mathbf{c}_i, \mathbf{c}_j) = \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) p(y|\mathbf{c}_i, \mathbf{c}_j). \quad (22)$$


Also, since the KL divergence is convex, we can have the following inequality:

$$\begin{aligned} D_{\text{KL}} \left[ \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) \hat{p}(y|z_i, z_j) \right] \\ \leq \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|z_i, z_j)]. \end{aligned} \quad (23)$$

If we use equations Eqs. (21) to (23) on the left-hand side of the inequality Eq. (20), we will have the following:

$$\begin{aligned} D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|\mathbf{x}_i, \mathbf{x}_j)] &= D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) \hat{p}(y|z_i, z_j)] \\ &\leq \sum_{z_i=1}^K \hat{p}(z_i|\mathbf{x}_i) \sum_{z_j=1}^K \hat{p}(z_j|\mathbf{x}_j) D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \parallel \hat{p}(y|z_i, z_j)]. \end{aligned}$$

Similar to supervised version for  $D_{\text{KL}}[p(y|\mathbf{c}_i, \mathbf{c}_j) \|\hat{p}(y|z_i, z_j)]$ , two possible scenarios can happen, i.e.,  $\mathbf{c}_i = \mathbf{c}_j$  and  $\mathbf{c}_i \neq \mathbf{c}_j$ . Fig. 4 depicts these different scenarios in which the model makes a mistake in assigning the label. We call these two scenarios intra-class divergence and inter-class divergence, respectively. Note that since the dataset is balanced, again we can consider that  $|z_i| = |z_j| = |\mathbf{c}_i| = |\mathbf{c}_j| = \frac{N}{K}$ . We also assume  $\hat{p}(\mathbf{c}_i) \approx p(\mathbf{c}_i) = \hat{p}(z_i) = \frac{1}{K}$ . In the transition from supervised to unsupervised scenario, the adaptation of the supervised derivations necessitates simply substituting  $\mathbf{c}_i$  with  $z_i$ . For brevity, we omit the derivation of these adapted equations.



**Fig. 4: Different scenarios for two sample category variables in unsupervised contrastive learning.** The latent variables  $z_i$  and  $z_j$  are extracted from the model. Model does the red assignments, while blue assignments indicate the ground truth category variable. For unsupervised contrastive learning, neither of the labels is known, so the model makes a mistake if (c)  $i$  and  $j$  have different labels and are assigned to the same label, (d)  $i$  and  $j$  have the same label, but model assign them to different labels.

## 2.5 Self-Expertise.

In the paper, we introduced an upper bound for a dataset comprising  $N$  samples across  $K$  categories, denoted as  $\mathcal{S}_K$ . To further refine our understanding, we propose an alternative upper bound,  $\hat{\mathcal{S}}_K$ , which leverages the hierarchical structure inherent among the categories. This hierarchical framework can be conceptualized as a Markov chain, where each category’s relevance is determined independently of the samples and hyper-labels, except for its immediate predecessor and its children in the hierarchy. Our objective is to delineate the conditions under which  $\hat{\mathcal{S}}_K$  serves as an effective upper bound, taking into account the hierarchical relationships among categories. This approach aims to capture the nuanced dependencies within the category structure, thereby providing a more granular and accurate upper limit for the distribution of samples among the categories.

$$\sum_{l=1}^{\lg K} D_{\text{KL}}[p(y|\mathbf{c}_i^l, \mathbf{c}_j^l) \|\hat{p}(y|\mathbf{c}_i^l, \mathbf{x}_j, \mathbf{c}_j^{l-1})] \leq \hat{\mathcal{S}}_K, \quad (24)$$

let  $l$  denote the hierarchy level. Consider that at level  $l$ , given  $K$  distinct categories, each cluster  $\mathbf{c}_j^{l-1}$  contains  $\frac{N}{K}$  samples, as opposed to the full  $N$  samples without considering the hierarchy. Upon substituting these adjusted values for  $K$  and  $\frac{K}{2}$  into the established upper bound in Eq. (4), we obtain the following:

$$\hat{\mathcal{S}}_K = \lg K \ln \frac{N}{K^2}, \quad \hat{\mathcal{S}}_{\frac{K}{2}} = (\lg K - 1) \ln \frac{4N}{K^2} \Rightarrow \hat{\mathcal{S}}_K \leq \hat{\mathcal{S}}_{\frac{K}{2}}. \quad \square \quad (25)$$

Employing analogous reasoning for the unsupervised self-expertise component, we derive:

$$\hat{\mathcal{U}}_K \leq \hat{\mathcal{U}}_{\frac{K}{2}}. \quad (26)$$

### 3 Method

For a deeper understanding of the proposed methods, this appendix section provides a detailed explanation and corresponding pseudocode for the algorithms briefly introduced in the main text.

#### 3.1 Balanced Semi-Supervised K-means

Our algorithm consists of three key stages:

1. Semi-Supervised K-means Centers Initialization,
2. Updating novel categories while balancing clusters,
3. Semi-supervised K-means for final assignment based on the refined centers.

---


#### Algorithm 1 Enhanced Semi-Supervised K-Means Initialization

---

- 1: **Input:** Labeled data points  $x_i^l$  with labels  $y_i^l$ , unlabeled data points  $x_j^u$ , cluster size threshold  $C$ , counts of labeled categories  $Y_c^l$  and unlabeled categories  $Y_c^u$
  - 2: **Output:** Refined initial positions for cluster centroids
  - 3: **Procedure Initialization:** Initialize centroids  $c_1, \dots, c_k$  as the centroids of labeled data  $x_i^l$  for each category  $y_i^l = c_1, \dots, c_k$ . Exclude the  $C$  nearest data points to each of these initial centroids to refine the cluster quality.
  - 4: **repeat**
  - 5:   Select a random data point  $x_j$  from the pool of unlabeled data.
  - 6:   Exclude the  $C$  nearest neighbors of  $x_j$  to ensure diverse centroid selection.
  - 7:   Incorporate  $x_j$  into the collection of centroids.
  - 8: **until** Achievement of  $K$  centroids
- 

**Semi-Supervised K-means Centers Initialization.** In our proposed method, Balanced Self-Supervised K-means (BSSK), the initial step involves the identification of K-means cluster centers for pre-labeled categories. This process is executed by calculating the cluster centers from the labeled data corresponding to the known categories. In scenarios involving balanced datasets characterized by a cluster size  $C$ , we enhance dataset purity by excluding the  $C$  nearest samples surrounding each data point to mitigate cluster overlap and ensure distinct cluster formation. Subsequently, for the incorporation of novel categories into the learning framework, BSSK adopts a strategy of randomly selecting samples from the pool of remaining data to serve as new cluster centers. Following the selection of each novel cluster center, we apply a similar purification step by removing the  $C$  closest samples in its vicinity. This systematic approach not only facilitates the exploration of new categories within the dataset but also maintains the integrity of cluster separation, ensuring a robust framework for the effective integration of both known and novel categories in self-supervised learning environments.

The pseudocode for this initialization is detailed in Algorithm 1, with a practical illustration provided in Fig. 5. It is important to highlight that in scenarios involving unbalanced datasets, this phase of the process is bypassed.



**Fig. 5: Cluster Center Initialization for Balanced Semi-Supervised K-Means.** Leveraging the advantage of having access to samples from known categories, we initiate the cluster centers for these categories directly using the available partial data from corresponding clusters (shown with green crosses). Given the predefined cluster size, denoted as  $C$ , it is prudent to avoid selecting initial centers for new categories from within the immediate vicinity of  $C$  of these established centers. Furthermore, upon determining a center for each novel category, we deliberately exclude its  $C$  closest samples in the dataset. This approach ensures that the initial centers for K-means clustering are adequately spaced apart, promoting a more balanced and effective clustering process.

Instead, we adopt an initial cluster center selection mechanism akin to that used in Semi-Supervised K-means [59]. This adjustment ensures our method remains robust and adaptable to diverse dataset characteristics.

---

#### Algorithm 2 Refinement of Novel Category Centers

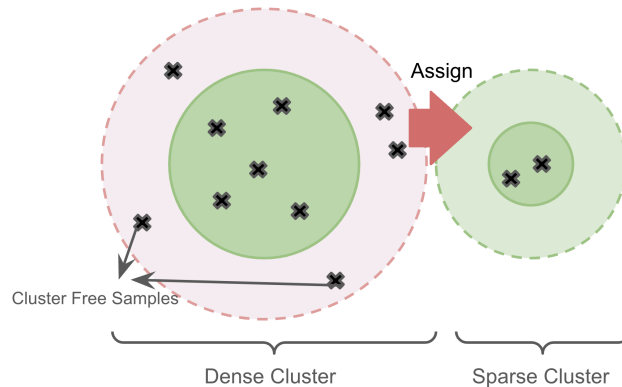
---

**Require:** Defined cluster size  $C$ , Initial cluster centers  $\mu^1, \dots, \mu^K$

**Ensure:** Generation of pseudo-labels for unlabeled data

- 1: **for** each iteration  $iter = 1$  to  $n_{iter}$  **do**
  - 2:   Assign each data point  $x_j$  to its nearest cluster center.
  - 3:   **if** specified cluster size  $C$  is not null **then**
  - 4:     Ensure each cluster adheres to the defined size  $C$  by applying balancing strategy from Algorithm 3.
  - 5:   **end if**
  - 6:   Recalculate the centers for the updated clusters.
  - 7: **end for**
- 

**Updating novel categories while balancing clusters.** In our proposed approach, we diverge from the conventional Semi-Supervised K-means [59] methodology, which updates all cluster centers simultaneously. Our premise is that the cluster centers corresponding to known data categories are already well-established and reliable. However, traditional K-means does not ensure the optimization of cluster centers for these known categories. To address this, our method selectively updates the cluster centers for novel categories while maintaining the integrity of the known category centers. This strategy ensures that, by the conclusion of this process, the centers for known categories are optimally positioned, whereas



**Fig. 6: Balancing Cluster Allocation Through Stable Matching.** Following the assignment of clusters, it is anticipated that certain clusters will emerge as *Dense*, possessing a number of samples exceeding the designated cluster size  $C$ , while others will manifest as *sparse*, characterized by a sample count below  $C$ . A proposed strategy involves permitting the centers of Dense clusters to select their  $C$  nearest samples, thereafter releasing any additional samples previously assigned to them as *Cluster-free*. These *Cluster-free* samples will then gravitate towards and select the nearest sparse cluster available to them, thereby balancing the allocation.

the centers for novel categories have been effectively adjusted. This technique is detailed in Algorithm 2.

For balanced datasets, our method enforces uniformity in cluster size, setting  $C = \frac{N}{K}$ , through a stable matching algorithm. Initially, each sample is allocated to its nearest cluster. Subsequently, if a cluster receives assignments exceeding  $C$ , it retains only the  $C$  closest samples, releasing the others. These unassigned samples are then considered "cluster-free." Concurrently, clusters with fewer than  $C$  assignments are identified as sparse. Each cluster-free sample is then allocated to the nearest sparse cluster. This procedure ensures stability within the algorithm, guaranteeing that no sample-cluster pair would prefer an alternative arrangement to their current one. The pseudocode for achieving balance is outlined in Algorithm 3, with the corresponding matching process illustrated in Figure Fig. 6. The robustness and stability of our matching algorithm are further validated by a subsequent theorem, reinforcing the effectiveness of our approach within the framework of this study.

**Theorem 3.** *The balancing stage of BSSK is a stable matching.*

*Proof.* Suppose, for the sake of contradiction, that our assumption is false. Therefore, we have two pairs of sample clusters,  $(x_i, c_k)$  and  $(x_j, c_l)$ , which would prefer to be paired with each other but are currently assigned differently using the BSSK algorithm. This implies that  $x_j$  is closer to  $c_k$  than  $x_i$  is, and since  $c_k$  retains its closest assigned samples, it follows that either:

1.  $x_j$  was never assigned to  $c_k$  using K-means.
2.  $x_j$  was assigned to  $c_k$  by K-means but was later re-assigned to another cluster.

---

**Algorithm 3** Balanced Distribution of Category Clusters

---

**Require:** Designated cluster size  $C$ , Initial set of clusters  $c^1, \dots, c^K$ , Preliminary cluster assignments.

**Ensure:** Enhanced pseudo-labels for untagged data

- 1: **for**  $k = 1$  **to**  $K$  **do**
  - 2:   Arrange samples within cluster  $c^k$  in ascending order based on proximity to the central point  $\mu^k$ .
  - 3:   Determine the total sample count within cluster  $|c^k|$ .
  - 4:   **if**  $|c^k| > C$  **then**
  - 5:     Retain the  $C$  samples closest to  $\mu^k$  within  $c^k$ , categorizing the remainder as *cluster-free*.
  - 6:   **else if**  $|c^k| < C$  **then**
  - 7:     Designate cluster  $c^k$  as *sparse*.
  - 8:   **end if**
  - 9: **end for**
  - 10: Redirect all *cluster-free* samples  $x_j$  towards the nearest *sparse* cluster.
- 

In the first scenario, for the pairing to be considered unstable,  $x_j$  would need to prefer  $c_k$  over  $c_l$ . However, if this were the case,  $x_j$  should have been assigned to  $c_k$  instead of  $c_l$ , leading to a contradiction.

In the second scenario, if  $x_j$  was re-assigned by BSSK, it suggests that  $c_k$  was identified as a dense cluster. Consequently,  $x_i$  cannot be re-assigned to it later as it is not part of a sparse cluster, and  $x_j$  was not among the  $C$  nearest neighbours of  $c_k$  while  $x_i$  was. This scenario implies that  $c_k$  prefers  $x_i$  over  $x_j$ , contradicting the initial claim of instability in the pairing. Therefore, by contradiction, the theorem is proven.  $\square$

---

**Algorithm 4** Balanced Semi-Supervised Kmeans (BSSK)

---

**Require:** Cluster size  $C$ , Initial centers  $\mu^1, \dots, \mu^K$

**Ensure:** Pseudo-labels for unlabeled data

- 1: Use algorithm Algorithm 2 to find good initial centers for novel categories.
  - 2: **for**  $iter = 1$  **to**  $n_{iter}$  **do**
  - 3:   Assign each data point  $x_j$  to the closest cluster center.
  - 4:   Update centers of all clusters.
  - 5: **end for**
- 

**Final assignment using semi-supervised Kmeans.** In the culminating phase of our Balanced Semi-Supervised K-means (BSSK) methodology, we implement a traditional Semi-Supervised K-means algorithm, leveraging both known and novel clusters identified in earlier steps. This approach is motivated by two primary considerations. Firstly, the initial identification of moderately optimal novel and known cluster centers reduces the likelihood of K-means converging to suboptimal cluster centers, thereby enhancing the optimization process for both known and

novel clusters. Secondly, it introduces a degree of flexibility regarding cluster size constraints. This adaptability ensures that even if the predefined cluster sizes are not precisely accurate or strictly enforced across the dataset, the model retains its ability to accurately assign samples. Algorithm 4 outlines the comprehensive framework of our balanced semi-supervised strategy, which is instrumental in generating pseudo-labels for the subsequent levels of category hierarchy.

---


**Algorithm 5** Hierarchical Semi-Supervised Kmeans (HSSK)
 

---

**Require:** Input data  $X$ , labeled samples  $\mathcal{Y}_S$ , initial cluster size  $C$

**Ensure:** Hierarchical Pseudo-labels for unlabeled data

- 1: **Initialization:** Apply BSSK to obtain initial pseudo-labels  $\hat{y}_j^1$  and cluster centers  $\mu_1^1, \dots, \mu_K^1$  for unlabeled data.
  - 2: Set hierarchy level  $h \leftarrow 1$ .
  - 3: **while** number of labeled prototypes  $n_i/2 > 1$  **do**
  - 4:   Perform K-means clustering of labeled prototypes into  $n_i/2$  groups.
  - 5:   Assign hyperlabels  $y^h$  to corresponding labeled data  $y$  using K-means results.
  - 6:   Update number of labeled prototypes:  $n_i \leftarrow n_i/2$ .
  - 7:   Update cluster size:  $C \leftarrow 2C$ .
  - 8:   Apply BSSK with updated  $C$  and hyperlabels  $y^h$  to obtain pseudo-labels and prototypes for the current level, resulting in  $\mu^h$  and  $\hat{y}^h$ .
  - 9: **end while**
- 



**Fig. 7: A single step of the Hierarchical Semi-Supervised K-Means (HSSK) algorithm.** Red circles represent novel category prototypes, and green circles denote known category prototypes, with the cluster centers for each indicated by a star symbol. Transitioning from level  $k$  to  $k+1$ , HSSK initially divides the prototypes of known categories in half, and applies the same partitioning process to the prototypes of novel categories. Subsequently, each sample within these clusters is assigned to the corresponding hyper label associated with the cluster centers. This methodology is anticipated to compel the model to segregate the representations of novel and known categories, thereby enhancing its discriminative capability.


### 3.2 Hierarchical Semi-Supervised K-means

In our methodology, the initial step involves the generation of pseudo-labels through the application of the Balanced Semi-Supervised K-means (BSSK) algo-

rithm. This foundational stage establishes our hierarchy’s base level, ensuring that the generated pseudo-labels correspond in granularity to the existing ground truth categories. We achieve this by determining  $K=|\mathcal{Y}_U|$  clusters and assigning them centers  $\mu_1^1, \dots, \mu_K^1$ , with the initial subset of these centers  $\mu_1^1 \dots \mu_{|\mathcal{Y}_S|}^1$  designated for the clusters associated with labeled data. The notation’s superindex 1 signifies this is the first level of abstraction.

Subsequently, for each subsequent level  $k$ , we conduct a twofold operation: halving the cluster centers of known prototypes from the previous level and applying the same halving procedure to cluster centers representing novel categories. We then reassign all samples linked to clusters at the  $(k-1)^{th}$  level to the newly formed clusters at level  $k$ , based on their hyperlabel associations. This iterative process is carried out until we delineate the most abstract dichotomy between ‘known’ and ‘novel’ categories. In scenarios involving balanced datasets, it is also pertinent to note that clusters at level  $k$  are designed to be double the size of those at level  $k-1$ , thus maintaining a consistent growth pattern across levels. For each iteration of the BSSK algorithm corresponding to a specific level  $k$ , we tailor the algorithm to accommodate the clustering characteristics of that level.

The procedural details of this Hierarchical Semi-Supervised K-means (HSSK) approach are succinctly described in our pseudocode, referenced in Algorithm 5. Additionally, we provide a visual representation of the HSSK mechanism, illustrated in Fig. 7, to show a single step of hierarchical structuring.



**Fig. 8: Construction Process of the Unsupervised Self-Expertise Target Matrix.** (a) Each sample is uniquely identified, ensuring distinctiveness. (b) Non-feline samples are associated with an uncertainty level of 0.5, allowing for moderate differentiation ambiguity. (c) Non animal samples are differentiated with an additional uncertainty of 0.25, ensuring more focus on distinguishing among animal categories. (d) The final aggregated matrix, subject to normalization, emphasizes high confidence in distinguishing between individual cats, while permitting greater uncertainty for non-animal entities, such as cars.

### 3.3 Unsupervised Self-Expertise

In the main text, we outlined the development of an unsupervised self-expertise target matrix designed to prioritize the model’s attention on the most challenging negative samples. The pseudocode for extracting these hard negatives via a Label


**Algorithm 6** Hierarchical Label Smoothing

**Require:** Unlabeled data samples  $x_j^u$ , desired cluster size  $C$ , and cluster centers  $\mu^1, \dots, \mu^K$

**Ensure:** Refined smoothed target matrix  $Y$ .

- 1: **Initialization:**
- 2: **for** each cluster  $k = 1$  to  $K$  **do**
- 3:     Determine cluster radius  $r^k$  for cluster  $k$
- 4: **end for**
- 5: Compute pairwise distances  $d_{ij}$  between data samples
- 6: Calculate label smoothing weights  $y_{ij}$  using:  $y_{ij} = \sum_k \mathbb{1}(d_{ij} > r^k) / 2^k$  and  $y_{ii} = 1$
- 7: Normalize the  $Y$  matrix to ensure validity.

Smoothing target matrix is detailed in Algorithm 6. Additionally, the incremental formation of this matrix is visually depicted in Fig. 8.



**Fig. 9: Construction Process of the Supervised Self-Expertise Target Matrix.**

(a) In the matrix, same-category pseudo-labels are treated as partially positive (0.5) with all others marked negative, targeting the first half of the latent representation. (b) Feline samples are assigned a positivity level of 0.25, aimed at enhancing their similarity detection, affecting only the first quarter of latent dimensions. (c) A distinction is made for all ‘known’ samples with an added uncertainty level of 0.125, concentrating on differentiating between known and unknown samples, impacting one-eighth of the latent space. (d) The comprehensive matrix, applied to the initial one-eighth of latent dimensions and normalized, focuses on enhancing category generalization.

**Intuition of Unsupervised Self-Expertise.** Considering visually similar samples within the same cluster as negatives can seem counter-intuitive at first. However, consider two visually similar yet semantically distinct samples, such as ‘sparrow’ and ‘finch,’ misclassified under the same category of ‘sparrow.’ Our method actively differentiates ‘finch’ from ‘sparrow’ despite their visual resemblance. Note that for fine-grained datasets with many categories, the majority of visually similar samples do not share the same semantic categories. By treating the visually similar samples as harder negatives, we make our unsupervised expertise focus on these samples while supervised self-expertise does the more coarse distinctions at the category level.

### 3.4 Supervised Self-Expertise

In alignment with the paradigm of unsupervised self-expertise, our approach necessitated the construction of a target matrix specific to supervised self-expertise, designed to encapsulate both weakly positive instances and strongly negative examples. Figure 9 elucidates the stepwise development of this matrix across various strata of the category hierarchy.

**Justification of Latent Utilization.** The rationale behind allocating the first half of the latent dimension to more abstract categories is twofold. Firstly, this allocation allows the model to make a rough categorization *e.g.* ‘feline’ or ‘canine’. Subsequently, the differentiation between precise categories *e.g.* ‘cat’ and ‘lion’ is handled by the second half. Utilizing the full dimension to simultaneously repel ‘cat’ and ‘lion’ samples, while attracting them at a higher abstraction level, impedes the model’s capacity to distinguish between these categories. This is primarily because the dominance of ‘feline’ samples obscures the finer distinctions between ‘cat’ and ‘lion’. By segregating ‘cat’ and ‘feline’ expertise, we ensure that a segment of the latent dimension is exclusively dedicated to discerning between finer-grained categories.

**Table 1: Statistics of datasets and their data splits for the generalized category discovery task.** The first four datasets are fine-grained image classification datasets, while the next three are coarse-grained datasets. The Herbarium19 dataset is both fine-grained and long-tailed.

DATASET	LABELLED		UNLABELLED	
	#IMAGES	#CATEGORIES	#IMAGES	#CATEGORIES
CUB-200 [61]	1.5K	100	4.5K	200
FGVC-AIRCRAFT [38]	3.4K	50	6.6K	100
STANFORD-CARS [32]	2.0K	98	6.1K	196
OXFORD-PET [47]	0.9K	19	2.7K	37
CIFAR-10 [33]	12.5K	5	37.5K	10
CIFAR-100 [33]	20.0K	80	30.0K	100
IMAGENET-100 [15]	31.9K	50	95.3K	100
HERBARIUM19 [57]	8.9K	341	25.4K	683

## 4 Experiments

### 4.1 Experimental Setup

**Datasets.** In this work, we evaluate the performance of our proposed method across several datasets to demonstrate its effectiveness and versatility. We conduct experiments on four detailed datasets, namely Caltech-UCSD Birds-200-2011 (CUB-200) [61], Fine-Grained Visual Classification of Aircraft (FGVC-Aircraft) [38], Stanford Cars [32], and Oxford-IIIT Pet [47], to assess its efficacy in fine-grained image recognition tasks. These tasks require the identification of subtle distinctions among highly similar categories, such as different bird species, aircraft models, car brands, and pet breeds, thereby posing significant challenges for image recognition systems. Furthermore, we extend our evaluation to more general coarse-grained datasets, including CIFAR10, CIFAR100 [33], and a subset of ImageNet comprising 100 categories, referred to as ImageNet-100 [15]. This extension showcases the adaptability of our method to broader classification challenges beyond fine-grained tasks. CIFAR10 and CIFAR100 feature coarse-grained categories encompassing a wide range of objects, such as vehicles and animals, while ImageNet-100 offers a diverse set of general categories from the larger ImageNet dataset. To provide a comprehensive overview of our experiments, we include a detailed summary of the datasets’ characteristics, including their statistical distribution and train/test splits, in Tab. 1.

**CUB-200** [61], or Caltech-UCSD Birds-200-2011, emphasizes the importance of discerning minute details across different bird species, underlining the complexity of fine-grained image recognition tasks.

**FGVC-Aircraft** [38] challenges image recognition models with its focus on aircraft, where variations in design significantly affect structural appearances, highlighting the dataset’s fine-grained nature.

**Stanford Cars** [32] introduces the complexity of recognizing car brands from various angles and in different colors, showcasing the nuanced challenges within fine-grained datasets.

**Oxford-IIIT Pet** [47] centers on the classification of cat and dog breeds, where limited data availability increases the risk of overfitting, stressing the dataset’s fine-grained and challenging aspects.

**CIFAR10/100** [33] are broad-spectrum datasets that encompass a variety of general object categories, including vehicles and animals, illustrating the diversity of challenges in coarse-grained image classification.

**ImageNet-100** [15] serves as a representative subset of the extensive ImageNet database, focusing on 100 varied categories to evaluate the generalizability of our approach in a broader context.

**Herbarium-19** [57], although not mentioned in the main text, could be considered in further discussions as it represents a unique dataset focused on plant species and is long-tailed, illustrating the application of fine-grained classification in when categories are unbalanced. Its inclusion could enrich the diversity of datasets evaluated in the paper, emphasizing the method’s applicability across different domains.

**Implementation Details.** In our experiments, we adhered to the dataset division proposed by Vaze *et al.* [59], where half of the categories in each dataset are designated as known, except for CIFAR100, where 80% are reused as known categories. The labeled set consists of 50% of the samples from these known categories. The remainder of the known category data, along with all data from novel categories, comprise the unlabeled set. Following [59]; we use ViT-B/16 as our backbone, which is either pre-trained by DINOv1 [8] on unlabelled ImageNet 1K [34], or pretrained by DINOv2 [42] on unlabelled ImageNet 22K. We use the batch size of 128 for training and set  $\lambda=0.35$ . For label smoothing, we use  $\alpha=0.5$  for fine-grained datasets and  $\alpha=0.1$  for coarse-grained datasets. Different from [59], we froze the first 10 blocks of ViT-B/16 and fine-tuned the last two blocks instead of only the last one to have more parameters given that for each level, only a fraction of the latent dimension is considered.<sup>4</sup>

**Evaluation Metrics.** We use our proposed Balanced Semi-Supervised K-means (BSSK) to cluster the extracted embeddings. Then, we use the Hungarian algorithm [66] to solve the optimal assignment of emerged clusters to their ground truth labels. We report the accuracy of the model’s predictions on *All*, *Known*, and *Novel* categories. Accuracy on *All* is calculated using the whole unlabelled test set, consisting of known and unknown categories. For *Known*, we consider the samples with labels known during training. Finally, for *Novel*, we consider samples from the unlabelled categories at train time.

## 4.2 Additional Empirical Results.

**Aggregated Performance on Different Kinds of Datasets.** In this section, we present a comparative analysis of our method against other approaches, as

<sup>4</sup> Our code is available at: <https://github.com/SarahRastegar/SelEx>

summarized in Table 2. The table consolidates the performance metrics across different datasets. It is evident that our method not only demonstrates respectable performance across both kinds of datasets but also shows superior performance in fine-grained datasets. Furthermore, when considering the aggregated performance metrics, it becomes clear that our method surpasses the competing methods, underscoring its effectiveness and robustness in a variety of settings.


**Table 2: Comparison with state-of-the-art for Aggregated datasets** Bold numbers show the best accuracies. Our method has a consistent performance for the three experimental settings (*All*, *Known*, *Novel*) on fine and coarse-grained datasets.

		Fine-grained			Coarse-grained			Average		
		All	Known	Novel	All	Known	Novel	All	Known	Novel
DINOv1	GPC [79]	44.5	51.7	39.9	80.4	91.9	71.3	62.5	71.8	55.6
	GCD [59]	45.1	51.8	41.8	79.5	88.0	73.7	62.3	69.9	57.8
	XCon [17]	46.8	52.5	44.0	82.6	90.7	75.1	64.7	71.6	59.6
	PromptCAL [76]	55.1	62.2	51.7	<b>87.4</b>	91.2	84.0	71.3	76.7	67.9
	SimGCD [65]	56.1	65.5	51.5	86.7	89.8	<u>84.6</u>	71.4	77.7	68.1
	GCA [43]	58.4	67.5	54.0	86.9	<u>91.9</u>	<u>82.7</u>	<u>72.7</u>	79.7	68.4
	AMEND [3]	58.0	70.2	52.0	87.0	89.1	<b>86.5</b>	72.5	79.7	<u>69.3</u>
	InfoSieve [49]	<u>60.5</u>	<b>72.1</b>	<u>54.7</u>	84.5	91.2	79.2	72.5	<u>81.7</u>	67.0
<b>SelEx (Ours)</b>		<b>63.0</b>	<u>71.9</u>	<b>58.8</b>	<u>87.1</u>	<b>92.3</b>	83.0	<b>75.1</b>	<b>82.1</b>	<b>70.9</b>

**Long-tailed Datasets.** In this section, we present Tab. 3 which outlines the efficacy of our proposed method when applied to the Herbarium dataset, known for its long-tailed distribution and fine-grained categorization. Numerous real-world datasets demonstrate a power-law distribution. By simply omitting our algorithm’s cluster size balancing component, we can apply it to long-tailed scenarios. Despite our model’s initial assumption of category uniformity, using this trick, the results depicted in the table reveal that it achieves respectable performance on this challenging dataset. This table shows enhancements over the baseline GCD method, with gains of 4.2% overall, 3.9% in known, and 4.3% in novel categories, surpassing the baseline method GCD. This observation suggests that self-expertise is advantageous not only for uniform category distributions but also for non-uniform ones.

**Table 3: Comparison with state-of-the-art for long-tailed image classification on Herbarium-19 dataset.** Bold and underlined numbers, respectively, show the best and second-best accuracies. Our method improves the baseline GCD [59] for the three experimental settings (*All*, *Known*, and *Novel*). This table shows that our method is also suited for long-tailed settings with unbalanced categories.

		All	Known	Novel
DINOv1	<b>Method</b>			
	GCD [59]	35.4	51.0	27.0
	CMS [14]	36.4	54.9	26.4
	InfoSieve [49]	41.0	55.4	33.2
	PIM [13]	42.3	56.1	34.8
	SPTNet [62]	43.4	58.7	35.2
	SimGCD [65]	44.0	58.0	<u>36.4</u>
	AMEND [3]	<u>44.2</u>	<u>60.5</u>	35.4
	$\mu$ GCD [60]	<b>45.8</b>	<b>61.9</b>	<b>37.2</b>
<b>SelEx (Ours)</b>		39.6	54.9	31.3




**Fig. 10: Impact of varying labeled sample ratios for known categories.** Illustration of the model’s proficiency in identifying novel categories across a spectrum of labeled data ratios for known categories. Showcases the model’s consistent performance in discovering novel categories, even at a 0% labeling ratio—akin to unsupervised clustering—versus the optimal 100% labeling scenario, which transforms the task into one of novel class discovery.

### 4.3 Ablative Studies

**Impact of varying labeled sample ratios for known categories.** In this section, we explore the impact of varying the proportion of labeled samples from known categories during training on model performance. In the primary experiments detailed in the main body of the text, labels were available for 50% of the samples from known categories. As illustrated in Fig. 10, we investigate the consequences of altering the ratio of labeled samples, ranging from nearly 0%—which essentially positions the model as an unsupervised clustering approach—to 100%, where the problem can be posed as novel class discovery. Notably, we observe an anticipated enhancement in performance for known categories as the volume of labeled samples increases. Intriguingly, the performance on novel categories remains robust, achieving a score of 68 with only a 20% labeling ratio, which is on par with the 50% ratio utilized in our main experiments. This finding indicates that our model is capable of identifying novel categories effectively, even with a significantly reduced number of labeled samples for known categories.

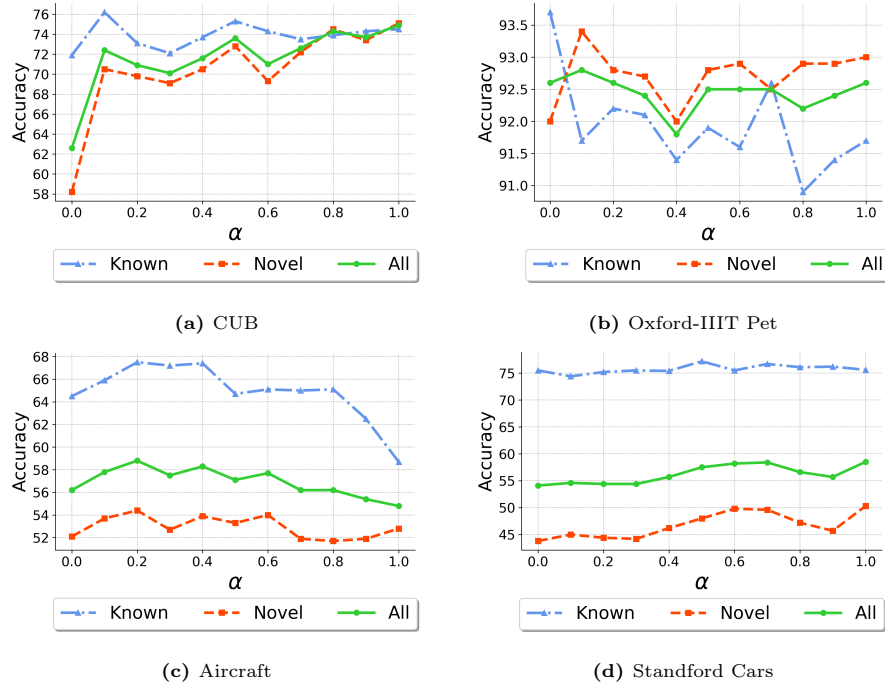
**Impact of the Supervised and Unsupervised Mixing Hyperparameter  $\lambda$ .** In this section, we modulate the mix between supervised learning and unsupervised self-expertise by varying the parameter  $\lambda$ , assessing its impact on model performance across different levels. The findings are summarized in Fig. 11. An overarching observation is that an increase in  $\lambda$  generally enhances performance, particularly in the context of novel category recognition. This improvement is attributed to our utilization of pseudo-labels for the supervision of novel categories, contrasting with the Generalized Category Discovery (GCD), which prioritizes optimization for known category classification, often at the expense of overall performance. The Self-Expertise (SelEx) approach, in contrast, leverages pseudo-labels for novel category supervision, leading to superior performance



**Fig. 11: Impact of varying  $\lambda$  values on performance.** highlighting the balance between supervised and unsupervised learning for novel category recognition. Demonstrates increased performance in novel categories with higher  $\lambda$ , and discusses the plateau in known category performance beyond a  $\lambda$  of 0.2, due to pseudo-label noise and diminished unsupervised learning effectiveness. This emphasizes the importance of optimizing  $\lambda$  to enhance novel category detection while managing the limitations posed by noise in pseudo-labels and unsupervised learning approaches.

in these categories as  $\lambda$  increases. However, beyond a  $\lambda$  value of 0.2, there is a negligible impact on the performance of known categories, with a slight decline observed. This can likely be attributed to the inherent noise in pseudo-labels, which adversely affects known category accuracy. Furthermore, with a higher  $\lambda$ , the effectiveness of unsupervised self-expertise diminishes, indicating a potential compromise in the model’s ability to differentiate between semantically distinct yet visually similar samples.


**Effects of smoothing hyperparameter** In this section, paralleling our main discussion, we delve into the impacts of the hyperparameter  $\alpha$  across various datasets, building upon our initial introduction of  $\alpha$  within the context of unsupervised self-expertise. To recap,  $\alpha$  serves as a smoothing hyperparameter that adjusts the uncertainty threshold for excluding negative samples from clusters other than the one under consideration. To elucidate, with  $\alpha=1$ , our model limits its incorporation of negative samples to those originating within its own cluster. On the flip side, an  $\alpha=0$  setting treats all negative samples uniformly, resonating with the principles of conventional unsupervised contrastive learning approaches. Through a series of experiments, whose  $\alpha$  values are meticulously outlined in Fig. 12, we observed a notable uplift in model efficacy on unseen categories as  $\alpha$  escalated. This boost is chiefly ascribed to the inherent limitation of traditional unsupervised contrastive learning methods, which indiscriminately push apart all non-identical samples, neglecting their potential semantic proximities. Our graphical representations underscore a general trend: as  $\alpha$  ascends, the model’s performance on novel categories improves, albeit at the expense of its effectiveness on familiar categories. Noteworthy is the pattern observed across datasets, except for Pets, where an initial performance increment for known categories is visible



**Fig. 12: Impact of Label Smoothing Hyperparameter  $\alpha$  Across Various Datasets.** This figure illustrates the differential effects of label smoothing, with increased smoothing typically enhancing performance on novel categories but impeding it on known ones. Furthermore, the benefits of more extensive smoothing are more pronounced for fine-grained datasets such as CUB, in contrast to Stanford Cars.

until reaching a saturation point. Beyond this plateau, the model’s performance either stabilizes—as observed in the CUB and Stanford Cars datasets—or begins to wane. This trend underscores the heightened efficacy of  $\alpha$  in amplifying performance for datasets with finer granularity, such as CUB, thereby suggesting the aptness of unsupervised self-expertise for more fine-grained datasets.

**Effects of number of frozen blocks** In this section, we delve into the impact of varying the quantity of frozen blocks on our model’s efficacy. It’s important to underline that an increase in the number of frozen blocks translates to a lesser deviation in the model’s weights from those of the pre-trained architecture. Specifically, a scenario with zero frozen blocks indicates complete parameter adaptability to the target dataset, whereas a setup with twelve frozen blocks signifies no alteration from the initial DINOv1 pre-trained weights. The outcomes, as illustrated in Figure Fig. 13, demonstrate a superior model performance with enhanced adaptability, particularly for recognizing novel categories. In contrast, for known categories, our findings reveal a roughly ascending trend, suggesting that minimal fine-tuning yields better results for these categories. Based on these observations, freezing eight blocks offers optimal results for novel categories, while,



**Fig. 13: Performance impact of varying frozen block counts on model adaptability.** This figure demonstrates the model’s enhanced ability to learn novel categories with fewer frozen blocks (optimal at 8 frozen blocks) and improved recognition of known categories with increased frozen blocks (optimal at 10 frozen blocks), highlighting the balance between adaptability and stability derived from the pretrained DINOv1.

as we have used in the main paper, freezing ten blocks is more advantageous for known categories.

**Table 4: Comparison with state-of-the-art with Estimated Number of Categories** Bold numbers show the best accuracies. We leverage the estimation technique from Vaze *et al.* [59] to determine category counts—231 for CUB and 230 for Stanford Cars. It’s noted that the category estimations by DCCL [48] and GCA [43] diverge from other methods. Our approach outperforms other methods in all test scenarios (*All*, *Known*, *Novel*), demonstrating robustness even with indeterminate category numbers.

Method	#Classes	CUB-200			Stanford-Cars			Average			
		All	Known	Novel	All	Known	Novel	All	Known	Novel	
DINOv1	GCD [59]	Known	51.3	56.6	48.7	39.0	57.6	29.9	45.2	57.1	39.3
	DCCL [48]	Known	63.5	60.8	64.9	43.1	55.7	36.2	53.3	58.3	50.6
	SimGCD [65]	Known	60.3	65.6	57.7	53.8	71.9	45.0	57.1	68.8	51.4
	GCA [43]	Known	<b>68.8</b>	<b>73.4</b>	<b>66.6</b>	54.4	<b>72.1</b>	45.8	<b>61.6</b>	<b>72.8</b>	56.2
	$\mu$ GCD [60]	Known	65.7	68.0	64.6	<u>56.5</u>	68.1	<b>50.9</b>	61.1	68.1	<u>57.8</u>
<b>SelEx (Ours)</b>	Known	<b>73.6</b>	<b>75.3</b>	<b>72.8</b>	<b>58.5</b>	<b>75.6</b>	<u>50.3</u>	<b>66.1</b>	<b>75.5</b>	<b>61.6</b>	
DINOv1	GCD [59]	Estimated	47.1	55.1	44.8	35.0	56.0	24.8	41.1	55.6	34.8
	DCCL [48]	Estimated	<u>63.5</u>	60.8	<u>64.9</u>	43.1	55.7	36.2	53.3	58.3	50.6
	SimGCD [65]	Estimated	61.5	<u>66.4</u>	59.1	49.1	65.1	41.3	55.3	65.8	50.2
	GCA [43]	Estimated	62.0	65.2	60.4	54.5	<u>71.2</u>	46.5	58.3	<u>68.2</u>	53.5
	$\mu$ GCD [60]	Estimated	62.0	60.3	62.8	<u>56.3</u>	66.8	<b>51.1</b>	<u>59.2</u>	63.6	<u>57.0</u>
	<b>SelEx (Ours)</b>	Estimated	<b>72.0</b>	<b>72.3</b>	<b>71.9</b>	<b>58.7</b>	<b>75.3</b>	<u>50.8</u>	<b>65.4</b>	<b>73.8</b>	<b>61.4</b>

**Estimate the number of categories** In the main text, we operated under the premise that the number of categories was predetermined. This assumption may not hold in practical scenarios. Therefore, in this segment, we examine the efficacy



Fig. 14: Failure cases of our approach

of our proposed SelEx model when the category count is deduced using a readily available technique, and we compare our results with several existing methods. For the estimation of category counts, we employ the method proposed by Vaze et al. [59], yielding 231 categories for the CUB dataset and 230 for the Stanford Cars dataset. It is important to highlight that the category count estimations provided by DCCL [48] and GCA [43] deviate from those obtained by other methodologies. As depicted in Table 4, our SelEx model surpasses competing approaches across all evaluation settings (*All*, *Known*, *Novel*), showcasing its adaptability in scenarios with uncertain category counts. This adaptability is further underscored by the fact that our model’s performance improves, particularly for novel categories within the Stanford Cars dataset. This improvement can be attributed to the model’s hierarchical nature, which does not strictly depend on an exact category count. Specifically, the model likely allocates additional categories to novel classes, given the higher estimated count of novel categories. This allocation effectively mitigates the general tendency of Generalized Category Discovery (GCD) methods to favor known categories, thereby enhancing performance for novel classes.

**Time-Complexity Analysis.** The time-complexity of our method is influenced by the variant of K-means in our semi-supervised approach. With GPU-based K-means, which loads the entire dataset’s embedding in memory, the complexity is  $\mathcal{O}(TK)$ , where  $K$  is the number of categories, and  $T$  is the iteration count. Practical timing of the GPU variant is provided in Tab. 5.

Table 5: Time and Computational complexity of SelEx

Method	Accuracy on CUB			Estimated Performance		
	All	Known	Novel	Epoch time	FLOPs	Params
GCD (Baseline)	51.3	56.6	48.7	61.57 s	19.98 G	37.25 M
SelEx	73.6	75.3	72.8	84.29 s	40.43 G	37.25 M

**Failure Cases.** Fig. 14 highlights several instances where our model encounters challenges. Specifically, the model has difficulty distinguishing known categories when parts of the bird are occluded or reflected. For novel categories, failures typically arise due to occlusion or camouflage. These observations indicate potential areas for future work to enhance the robustness of our approach.

## 5 Discussion

### 5.1 Limitations

A significant constraint of methodologies grounded in hierarchical structures lies in their inherent assumption of dataset hierarchies. In scenarios where the dataset lacks clear hierarchical organization, this presupposition could be detrimental. While our method, SelEx, endeavors to mitigate this limitation by dynamically allocating portions of the latent dimension to represent higher hierarchical levels, it does not completely circumvent the issue. Furthermore, SelEx presupposes a balanced data distribution, a presumption that becomes particularly challenging in the context of datasets with a long-tailed distribution, such as the Herbarium dataset. This limitation highlights a critical area for further improvement and adaptation in our approach.

### 5.2 Future Works

Our work introduces SelEx as a novel approach for generalized category discovery, demonstrating promising results. The ablation studies reveal SelEx’s potential applicability beyond its initial scope, particularly in few-shot and low-shot learning scenarios. Additionally, its hierarchical structure suggests utility in tasks with partially labeled data, offering a versatile framework for handling such complexities. Moreover, by incorporating label smoothing techniques to introduce uncertainty, SelEx shows resilience against noisy labels, enhancing its robustness and reliability in real-world applications. These insights indicate that SelEx possesses broad applicability across various domains within computer vision, especially in environments characterized by limited or sparse supervision. Future research will explore the extension of SelEx to a wider array of computer vision challenges, capitalizing on its adaptability and effectiveness in scenarios where supervision is minimal or data labels are imperfect.

## References

1. An, W., Shi, W., Tian, F., Lin, H., Wang, Q., Wu, Y., Cai, M., Wang, L., Chen, Y., Zhu, H., et al.: Generalized category discovery with large language models in the loop. arXiv preprint arXiv:2312.10897 (2023) [3](#)
2. An, W., Tian, F., Zheng, Q., Ding, W., Wang, Q., Chen, P.: Generalized category discovery with decoupled prototypical network. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37 (2023) [2](#)
3. Banerjee, A., Kallooriyakath, L.S., Biswas, S.: Amend: Adaptive margin and expanded neighborhood for efficient generalized category discovery. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2101–2110 (2024) [2](#), [3](#), [20](#)
4. Bendale, A., Boulton, T.: Towards open world recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1893–1902 (2015) [2](#)
5. Bendale, A., Boulton, T.E.: Towards open set deep networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1563–1572 (2016) [2](#)
6. Cao, K., Brbic, M., Leskovec, J.: Open-world semi-supervised learning. In: Proceedings of the International Conference on Learning Representations (2022) [2](#)
7. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems* **33**, 9912–9924 (2020) [1](#)
8. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9650–9660 (2021) [1](#), [19](#)
9. Chapelle, O., Scholkopf, B., Zien, A.: Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks* **20**(3), 542–542 (2009) [2](#)
10. Chen, G., Peng, P., Wang, X., Tian, Y.: Adversarial reciprocal points learning for open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(11), 8065–8081 (2021) [2](#)
11. Chen, G., Qiao, L., Shi, Y., Peng, P., Li, J., Huang, T., Pu, S., Tian, Y.: Learning open set network with discriminative reciprocal points. In: European Conference on Computer Vision. pp. 507–522. Springer (2020) [2](#)
12. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International Conference on Machine Learning. pp. 1597–1607. PMLR (2020) [1](#)
13. Chiaroni, F., Dolz, J., Masud, Z.I., Mitiche, A., Ben Ayed, I.: Parametric information maximization for generalized category discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1729–1739 (2023) [2](#), [20](#)
14. Choi, S., Kang, D., Cho, M.: Contrastive mean-shift learning for generalized category discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024) [2](#), [20](#)
15. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009) [18](#), [19](#)
16. Du, R., Chang, D., Liang, K., Hospedales, T., Song, Y.Z., Ma, Z.: On-the-fly category discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11691–11700 (2023) [2](#)

17. Fei, Y., Zhao, Z., Yang, S., Zhao, B.: Xcon: Learning with experts for fine-grained category discovery. In: British Machine Vision Conference (2022) [20](#)
18. Fini, E., Sangineto, E., Lathuilière, S., Zhong, Z., Nabi, M., Ricci, E.: A unified objective for novel class discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9284–9292 (2021) [2](#)
19. Gao, F., Zhong, W., Cao, Z., Peng, X., Li, Z.: Opengcd: Assisting open world recognition with generalized category discovery. arXiv preprint arXiv:2308.06926 (2023) [2](#)
20. Ge, Z., Demyanov, S., Chen, Z., Garnavi, R.: Generative openmax for multi-class open set classification. arXiv preprint arXiv:1707.07418 (2017) [2](#)
21. Guo, Y., Xu, M., Li, J., Ni, B., Zhu, X., Sun, Z., Xu, Y.: Hcsc: Hierarchical contrastive selective coding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9706–9715 (2022) [3](#)
22. Han, J., Luo, P., Wang, X.: Deep self-learning from noisy labels. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5138–5147 (2019) [2](#)
23. Han, K., Rebuffi, S.A., Ehrhardt, S., Vedaldi, A., Zisserman, A.: Automatically discovering and learning new visual categories with ranking statistics. In: Proceedings of the International Conference on Learning Representations (2020) [2](#)
24. Han, K., Vedaldi, A., Zisserman, A.: Learning to discover novel visual categories via deep transfer clustering. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8401–8409 (2019) [2](#)
25. Hao, S., Han, K., Wong, K.Y.K.: Cipr: An efficient framework with cross-instance positive relations for generalized category discovery. arXiv preprint arXiv:2304.06928 (2023) [2](#)
26. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020) [1](#)
27. Hsu, Y.C., Lv, Z., Kira, Z.: Learning to cluster in order to transfer across domains and tasks. arXiv preprint arXiv:1711.10125 (2017) [2](#)
28. Hsu, Y.C., Lv, Z., Schlosser, J., Odom, P., Kira, Z.: Multi-class classification without multi-class labels. arXiv preprint arXiv:1901.00544 (2019) [2](#)
29. Kim, H., Suh, S., Kim, D., Jeong, D., Cho, H., Kim, J.: Proxy anchor-based unsupervised learning for continuous generalized category discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16688–16697 (2023) [2](#)
30. Kim, S., Kim, D., Cho, M., Kwak, S.: Proxy anchor loss for deep metric learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3238–3247 (2020) [3](#)
31. Kong, S., Ramanan, D.: Opengan: Open-set recognition via open data generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 813–822 (2021) [2](#)
32. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 554–561 (2013) [18](#), [19](#)
33. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. Rep. 0, University of Toronto, Toronto, Ontario (2009), <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf> [18](#), [19](#)
34. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems. vol. 25 (2012) [19](#)

35. Lang, N., Snæbjarnarson, V., Cole, E., Mac Aodha, O., Igel, C., Belongie, S.: From coarse to fine-grained open-set recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17804–17814 (2024) [3](#)
36. Liu, M., Roy, S., Zhong, Z., Sebe, N., Ricci, E.: Large-scale pre-trained models are surprisingly strong in incremental novel class discovery. arXiv preprint arXiv:2303.15975 (2023) [2](#)
37. Liu, Y., Cai, Y., Jia, Q., Qiu, B., Wang, W., Pu, N.: Novel class discovery for ultra-fine-grained visual categorization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17679–17688 (2024) [2](#)
38. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151 (2013) [18](#)
39. Neal, L., Olson, M., Fern, X., Wong, W.K., Li, F.: Open set learning with counterfactual images. In: European Conference on Computer Vision. pp. 613–628 (2018) [2](#)
40. Oliver, A., Odena, A., Raffel, C.A., Cubuk, E.D., Goodfellow, I.: Realistic evaluation of deep semi-supervised learning algorithms. In: Advances in Neural Information Processing Systems. vol. 31 (2018) [2](#)
41. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018) [6](#)
42. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: DINOv2: Learning robust visual features without supervision. Transactions on Machine Learning Research (2024) [1](#), [19](#)
43. Otholt, J., Meinel, C., Yang, H.: Guided cluster aggregation: A hierarchical approach to generalized category discovery. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2618–2627 (2024) [2](#), [3](#), [20](#), [24](#), [25](#)
44. Ouali, Y., Hudelot, C., Tami, M.: An overview of deep semi-supervised learning. arXiv preprint arXiv:2006.05278 (2020) [2](#)
45. Ouldnooghi, R., Kuo, C.W., Kira, Z.: Clip-gcd: Simple language guided generalized category discovery. arXiv preprint arXiv:2305.10420 (2023) [3](#)
46. Oza, P., Patel, V.M.: C2ae: Class conditioned auto-encoder for open-set recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2307–2316 (2019) [2](#)
47. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3498–3505 (2012) [18](#), [19](#)
48. Pu, N., Zhong, Z., Sebe, N.: Dynamic conceptual contrastive learning for generalized category discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023) [2](#), [24](#), [25](#)
49. Rastegar, S., Doughty, H., Snoek, C.G.M.: Learn to categorize or categorize to learn? self-coding for generalized category discovery. In: Advances in Neural Information Processing Systems (2023) [2](#), [3](#), [20](#)
50. Rastegar, S., Soleymani, M., Rabiee, H.R., Shojaee, S.M.: Mdl-cw: A multimodal deep learning framework with cross weights. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2601–2609 (2016) [3](#)

51. Rebuffi, S.A., Ehrhardt, S., Han, K., Vedaldi, A., Zisserman, A.: Semi-supervised learning with scarce annotations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition workshops. pp. 762–763 (2020) [2](#)
52. Rizve, M.N., Kardan, N., Shah, M.: Towards realistic semi-supervised learning. In: European Conference on Computer Vision. pp. 437–455. Springer (2022) [2](#)
53. Roy, S., Liu, M., Zhong, Z., Sebe, N., Ricci, E.: Class-incremental novel class discovery. In: European Conference on Computer Vision. pp. 317–333. Springer (2022) [2](#)
54. Scheirer, W.J., Jain, L.P., Boulton, T.E.: Probability models for open set recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **36**(11), 2317–2324 (2014) [2](#)
55. Scheirer, W.J., de Rezende Rocha, A., Sapkota, A., Boulton, T.E.: Toward open set recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **35**(7), 1757–1772 (2012) [1](#)
56. Shu, Y., Shi, Y., Wang, Y., Huang, T., Tian, Y.: P-odn: Prototype-based open deep network for open set recognition. Scientific Reports **10**(1), 7146 (2020) [2](#)
57. Tan, K.C., Liu, Y., Ambrose, B., Tulig, M., Belongie, S.: The herbarium challenge 2019 dataset. arXiv preprint arXiv:1906.05372 (2019) [18](#), [19](#)
58. Tan, Z., Zhang, C., Yang, X., Sun, J., Huang, K.: Revisiting mutual information maximization for generalized category discovery. arXiv preprint arXiv:2405.20711 (2024) [2](#)
59. Vaze, S., Han, K., Vedaldi, A., Zisserman, A.: Generalized category discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022) [2](#), [11](#), [19](#), [20](#), [24](#), [25](#)
60. Vaze, S., Vedaldi, A., Zisserman, A.: No representation rules them all in category discovery. Advances in Neural Information Processing Systems 37 (2023) [3](#), [20](#), [24](#)
61. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset (Jul 2011) [18](#)
62. Wang, H., Vaze, S., Han, K.: SPTNet: An efficient alternative framework for generalized category discovery with spatial prompt tuning. In: Proceedings of the International Conference on Learning Representations (2024) [3](#), [20](#)
63. Wang, Y., Wang, Y., Wu, Y., Zhao, B., Qian, X.: Beyond known clusters: Probe new prototypes for efficient generalized class discovery. arXiv preprint arXiv:2404.08995 (2024) [2](#)
64. Wang, Y., Zhong, Z., Qiao, P., Cheng, X., Zheng, X., Liu, C., Sebe, N., Ji, R., Chen, J.: Discover and align taxonomic context priors for open-world semi-supervised learning. In: Advances in Neural Information Processing Systems (2023) [3](#)
65. Wen, X., Zhao, B., Qi, X.: Parametric classification for generalized category discovery: A baseline study. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16590–16600 (2023) [2](#), [3](#), [20](#), [24](#)
66. Wright, M.: Speeding up the Hungarian algorithm. Computers & Operations Research **17**(1), 95–96 (1990) [19](#)
67. Wu, T.Y., Ho, C.H., Vasconcelos, N.: Protect: Prompt tuning for taxonomic open set classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16531–16540 (2024) [3](#)
68. Xiao, R., Feng, L., Tang, K., Zhao, J., Li, Y., Chen, G., Wang, H.: Targeted representation alignment for open-world semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23072–23082 (June 2024) [2](#)

69. Yang, M., Wang, L., Deng, C., Zhang, H.: Bootstrap your own prior: Towards distribution-agnostic novel class discovery. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3459–3468 (2023) [2](#)
70. Yang, X., Song, Z., King, I., Xu, Z.: A survey on deep semi-supervised learning. *IEEE Transactions on Knowledge and Data Engineering* (2022) [2](#)
71. Yang, Z., Bastan, M., Zhu, X., Gray, D., Samaras, D.: Hierarchical proxy-based loss for deep metric learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1859–1868 (2022) [3](#)
72. Yoshihashi, R., Shao, W., Kawakami, R., You, S., Iida, M., Naemura, T.: Classification-reconstruction learning for open-set recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4016–4025 (2019) [2](#)
73. Yu, Y., Wang, S., He, Y., Chen, J., Gong, Y.: Continual novel class discovery via feature enhancement and adaptation. *arXiv preprint arXiv:2405.06389* (2024) [2](#)
74. Yue, Z., Wang, T., Sun, Q., Hua, X.S., Zhang, H.: Counterfactual zero-shot and open-set visual recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15404–15414 (2021) [2](#)
75. Zhang, L., Qi, L., Yang, X., Qiao, H., Yang, M.H., Liu, Z.: Automatically discovering novel visual categories with self-supervised prototype learning. *arXiv preprint arXiv:2208.00979* (2022) [2](#)
76. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023) [2](#), [3](#), [20](#)
77. Zhang, S., Xu, R., Xiong, C., Ramaiah, C.: Use all the labels: A hierarchical multi-label contrastive learning framework. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16660–16669 (2022) [3](#)
78. Zhao, B., Han, K.: Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. *Advances in Neural Information Processing Systems* **34**, 22982–22994 (2021) [2](#)
79. Zhao, B., Wen, X., Han, K.: Learning semi-supervised gaussian mixture models for generalized category discovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2023) [2](#), [20](#)
80. Zheng, H., Pu, N., Li, W., Sebe, N., Zhong, Z.: Textual knowledge matters: Cross-modality co-teaching for generalized visual class discovery. *arXiv preprint arXiv:2403.07369* (2024) [3](#)
81. Zhong, Z., Zhu, L., Luo, Z., Li, S., Yang, Y., Sebe, N.: Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9462–9470 (2021) [2](#)