

Parametric Classification for Generalized Category Discovery: A Baseline Study

Xin Wen^{1*} Bingchen Zhao^{2*} Xiaojuan Qi¹
¹The University of Hong Kong ²University of Edinburgh
 {wenxin, xjq}@eee.hku.hk bingchen.zhao@ed.ac.uk

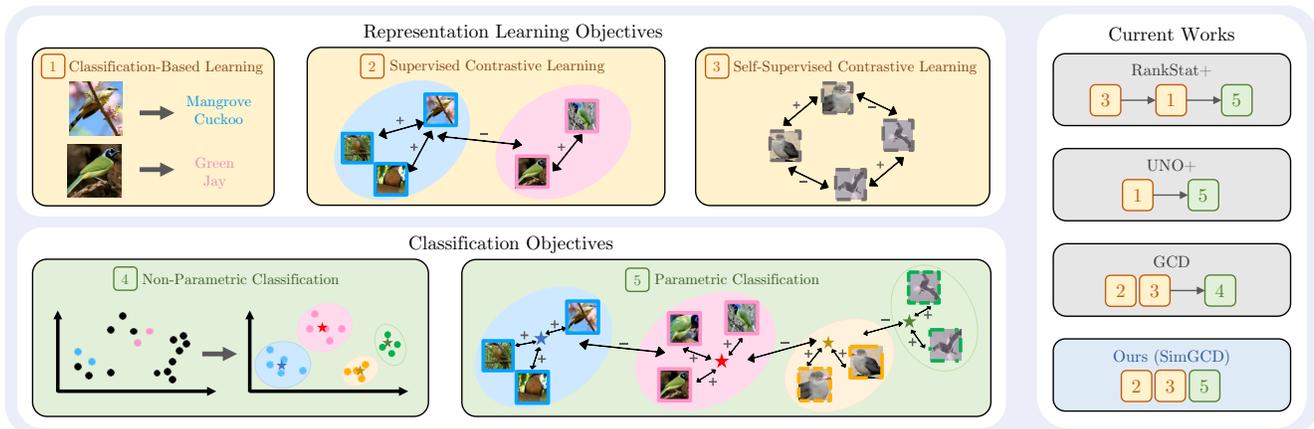


Figure 1. Left: building blocks for representation learning or classifier learning; Right: overall abstraction of current works, where ‘→’ separates different stages of the method. Our work builds on GCD [39], and jointly trains a parametric classifier.

Abstract

Generalized Category Discovery (GCD) aims to discover novel categories in unlabelled datasets using knowledge learned from labelled samples. Previous studies argued that parametric classifiers are prone to overfitting to seen categories, and endorsed using a non-parametric classifier formed with semi-supervised k -means. However, in this study, we investigate the failure of parametric classifiers, verify the effectiveness of previous design choices when high-quality supervision is available, and identify unreliable pseudo-labels as a key problem. We demonstrate that two prediction biases exist: the classifier tends to predict seen classes more often, and produces an imbalanced distribution across seen and novel categories. Based on these findings, we propose a simple yet effective parametric classification method that benefits from entropy regularisation, achieves state-of-the-art performance on multiple GCD benchmarks and shows strong robustness to unknown class numbers. We hope the investigation and proposed simple framework can serve as a strong baseline to facilitate future studies in this field. Our code is available at: <https://github.com/CVMI-Lab/SimGCD>.

*Equal contribution.

1. Introduction

With large-scale labelled datasets, deep learning methods can surpass humans in recognising images [23]. However, it is not always possible to collect large-scale human annotations for training deep learning models. Therefore, there is a rich body of recognition models that focus on learning with a large number of unlabelled data. Among them, semi-supervised learning (SSL) [29, 4, 33] is regarded as a promising approach, yet with the assumption that labelled instances are provided for each of the categories the model needs to classify. Generalized category discovery (GCD) [39] is recently formalised to relax this assumption by assuming the unlabelled data can also contain similar yet distinct categories from the labelled data. The goal of GCD is to learn a model that is able to classify the already-seen categories in the labelled data, and more importantly, *jointly discover the new categories in the unlabelled data and make correct classifications*. Developing a strong method for this problem could help us better utilise the easily available large-scale unlabelled datasets.

Previous works [39, 20, 16, 5] approach this problem from two perspectives: learning generic feature representations to facilitate the discovery of novel categories, and gen-

erating pseudo clusters/labels for unlabelled data to guide the learning of a classifier. The former is often achieved by using self-supervised learning methods [20, 46, 17, 22, 8, 48] to improve the generalization ability of features to novel categories. For constructing the classifier, earlier works [20, 46, 51, 5, 16] adopt a parametric approach that builds a learnable classifier on top of the extracted features. The classifier is jointly optimised with the backbone using labelled data and pseudo-labelled data.

However, recent research shows [39, 15] that parametric classifiers are prone to overfit to seen categories (see Fig. 2) and thus promote using a non-parametric classifier such as k -means clustering. Albeit obtaining promising results, the non-parametric classifiers suffer from heavy computation costs on large-scale datasets due to quadratic complexity of the clustering algorithm. Besides, unlike a learnable parametric classifier, the non-parametric method loses the ability to jointly optimise the separating hyperplane of all categories in a learnable manner, potentially being sub-optimal.

This motivates us to revisit the reason that makes previous parametric classifiers fail to recognise novel classes. In a series of investigations (Sec. 3) from the view of supervision quality, we verify the effectiveness of prior design choices in feature representations and training paradigms when strong supervision is available, and conclude that the key to previous parametric classifiers’ degraded performance is unreliable pseudo labels. By diagnosing the statistics of its predictions, we identify severe prediction biases within the model, *i.e.*, the bias towards predicting more ‘Old’ classes than ‘New’ classes (Fig. 5) and the bias of producing imbalanced pseudo-labels across all classes (Fig. 6).

Based on these findings, we thus present a simple parametric classification baseline for generalized category discovery (see Figs. 1 and 7). The representation learning objective follows GCD [39], and the classification objective is simply cross-entropy for labelled samples and self-distillation [8, 2] for unlabelled samples. Besides, an entropy regularisation term is also adopted to overcome biased predictions by enforcing the model to predict more uniformly distributed labels across all possible categories. Empirically, we indeed observe that our method produces more balanced pseudo-labels (Figs. 9 and 10) and achieves a large performance gain on multiple GCD benchmarks (Tabs. 2 to 4), indicating that the two types of biases we identified are the core reason why the parametric-classifier-based approach performs poorly for GCD. Additionally, we observe that the entropy regulariser could also be used to enforce robustness towards an unknown number of categories (Figs. 11 and 12), this could further ease the deployment of parametric classifiers for GCD in real-world scenarios.

Our contributions are summarised as follows: (1) We revisit the design choices of parametric classification and conclude the key factors that make it fail for GCD. (2) Based on

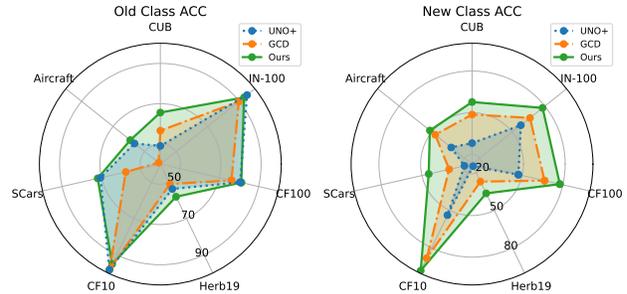


Figure 2. **Performance overview.** Prior parametric classification method (UNO+ [16]) shows highly degraded performance in ‘New’ classes. The non-parametric classification work (GCD [39]) performs better, but at the sacrifice of ‘Old’ class and high inference cost. Our method shows that parametric classification can work well on both metrics.

the analysis, we propose a simple yet effective parametric classification method. (3) Our method achieves SOTA on multiple popular GCD benchmarks, challenging the recent promotion of non-parametric classification for this task.

2. Related Works

Semi-Supervised Learning (SSL) has been an important research topic where a number of methods have been proposed [4, 33, 36]. SSL assumes that the labelled instances are available for all possible categories in the unlabelled dataset; the objective is to learn a model to perform classification using both the labelled samples as well as the large-scale available unlabelled data. One of the most effective methods for SSL is the consistency-based method, where the model is forced to learn consistent representations of two different augmentations of the same image [33, 4, 36]. Furthermore, it is also shown that self-supervised representation learning is helpful for the task of SSL [45, 30] as it can provide a strong representation for the task.

Open-Set Semi-Supervised Learning considers the case where the unlabelled data may contain outlier data points that do not belong to any of the categories in the labelled training set. The goal is to learn a classifier for the labelled categories from a noisy unlabelled dataset [44, 32, 10, 19]. As this problem only focuses on the performance of the labelled categories, the outlier from novel categories are simply rejected and no further classification is needed.

Generalized Category Discovery (GCD) is a relatively new problem recently formalised in Vaze *et al.* [39], and is also studied in a parallel line of work termed open-world semi-supervised learning [5, 34]. Different from the common assumption of SSL [29], GCD does not assume the unlabelled dataset comes from the same class set as the labelled dataset, posing a greater challenge for designing an effective model. GCD can be seen as a natural extension of the novel category discovery (NCD) problem [21]

where it is assumed that the unlabelled dataset and the labelled dataset do not have any class overlap, thus baselines for NCD [20, 46, 51, 50, 16] can be adopted for the GCD problem by extending the classification head to have more outputs [39]. The incremental setting of GCD is also explored [47, 31]. It is pointed out in [39] that a non-parametric classifier formed using semi-supervised k -means can outperform strong parametric classification baselines from NCD [20, 16] because it can alleviate the overfitting to seen categories in the labelled set. In this paper, we revisit this claim and show that parametric classifiers can reach stronger performance than non-parametric classifiers.

Deep Clustering aims at learning a set of semantic prototypes from unlabelled images with deep neural networks. Considering that no label information is available, the focus is on how to obtain reliable pseudo-labels. While early attempts rely on hard labels produced by k -means [6], there has been a shift towards soft labels produced by optimal transport [1, 7], and more recently sharpened predictions from an exponential moving average-updated teacher model [8, 2]. Deep clustering has shown strong potential for unsupervised representation learning [6, 1, 7, 8, 2], unsupervised semantic segmentation [11, 43], semi-supervised learning [3], and novel category discovery [16]. In this work, we study the techniques that make strong parametric classifiers for GCD with inspirations from deep clustering.

3. On the Failure of Parametric Classification

In order to explore the reason that makes previous parametric classifiers fail to recognise ‘New’ classes for generalized category discovery, this section presents preliminary studies to reveal the role of two major components: representation learning (Sec. 3.2) and pseudo-label quality on unseen classes (Sec. 3.3). These have led to *conflicting choices* of previous works, but why? We show a unified viewpoint (Figs. 3 and 4), and emphasise that *taking pseudo-label quality into account* is important for selecting the suitable design choice. This then led to our diagnosis of what makes the degenerated pseudo-labels (Sec. 3.4), and motivated our de-biased pseudo-labelling strategy.

3.1. Investigation Setting

Generalized category discovery. Given an unlabelled dataset $\mathcal{D}^u = \{(\mathbf{x}_i^u, y_i^u)\} \in \mathcal{X} \times \mathcal{Y}_u$ where \mathcal{Y}_u is the label space of the unlabelled samples, the goal of GCD is to learn a model to categorise the samples in \mathcal{D}^u using the knowledge from a labelled dataset $\mathcal{D}^l = \{(\mathbf{x}_i^l, y_i^l)\} \in \mathcal{X} \times \mathcal{Y}_l$ where \mathcal{Y}_l is the label space of labelled samples and $\mathcal{Y}_l \subset \mathcal{Y}_u$. We denote the number of categories in \mathcal{Y}_u as $K_u = |\mathcal{Y}_u|$, it is common to assume the number of categories is known *a-priori* [20, 46, 51, 16], or can be estimated using off-the-shelf methods [21, 39].

Representation learning. For representation learning, we follow GCD [39], which applies supervised contrastive learning [24] on labelled samples, and self-supervised contrastive learning [9] on all samples (detailed in Sec. 4.1).

Classifier. We follow UNO [16] to adopt a *prototypical classifier*. Take $f(\mathbf{x})$ as the feature vector of an image \mathbf{x} extracted using from the backbone f , the procedure for producing logits is $\mathbf{l} = \frac{1}{\tau}(\mathbf{w}/\|\mathbf{w}\|)^\top(f(\mathbf{x})/\|f(\mathbf{x})\|)$. Here τ is the temperature value that scales up the norm of \mathbf{l} and facilitates optimisation of the cross-entropy loss [42].

Training settings. We train with varying supervision qualities. The *minimal supervision* setting utilises only the labels in \mathcal{D}^l , while the *oracle supervision* setting assumes all samples are labelled (both \mathcal{D}^l and \mathcal{D}^u). Besides, we study two practical settings that adopt pseudo labels for unlabelled samples in \mathcal{D}^u : *self-label* that predicts pseudo-labels with the Sinkhorn Knopp algorithm following [16], and *self-distil*, which depicts another pseudo-labelling strategy as in Fig. 7 and will be introduced in detail in Sec. 4.2. For all settings, we only employ a cross-entropy loss on the (pseudo-)labelled samples on hand for classification. Note that unless otherwise stated, *this is done on decoupled features, thus representation learning is unaffected*.

3.2. Which Representation to Build Your Classifier?

Motivation. Following the trend of deep clustering that focuses on self-supervised representation learning [7], previous parametric classification work UNO [16] fed the classifier with representations taken from the *projector*. While in GCD [39], significantly stronger performance is achieved with a non-parametric classifier built upon representations taken from the *backbone*. We revisit this choice as follows.

Setting. Consider f as the feature backbone, and g is a multi-layer perceptron (MLP) projection head. Given an input image \mathbf{x}_i , the representation from the *backbone* can be written as $f(\mathbf{x}_i)$, and that from the *projector* is $g(f(\mathbf{x}_i))$.

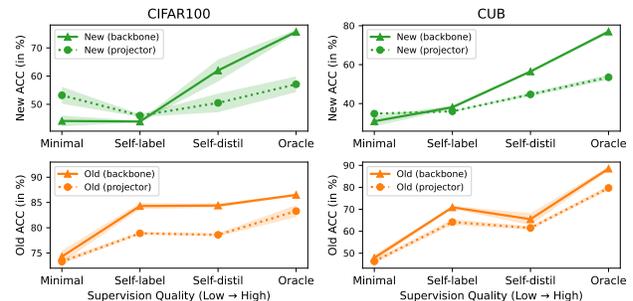


Figure 3. **Results with different representations.** We build the classifier on *post-backbone* or *post-projector* representations, and train with varying supervision quality. Results on ‘Old’ class consistently benefit from the post-backbone representations regardless of the supervision quality, while unleashing its potential on ‘New’ class requires stronger pseudo labels.

Result & discussion. As in Fig. 3, the post-backbone feature space has a clearly higher upper bound for learning prototypical classifiers than the post-projector feature space. Using a projector in self-supervised learning lets the projector focus on solving pretext tasks and allows the backbone to keep as much information as possible (which facilitates downstream tasks) [12]. But when good classification performance is all you need, our results suggest that the classification objective should build on post-backbone representations directly. The features post the projector might focus more on solving the pretext task and not be necessarily useful for the classification objective. Note that high-quality pseudo labels are necessary to unleash the post-backbone representations’ potential to recognise novel categories.

3.3. Decoupled or Joint Representation Learning?

Motivation. Previous parametric classification methods, e.g., UNO [16], commonly tune the representations jointly with the classification objective. On the contrary, in the two-stage non-parametric method GCD [39] where the performance in ‘New’ classes is notably higher, classification/clustering is fully decoupled from representation learning, and the representations can be viewed as unaltered by classification. In this part, we study whether the *joint* learning strategy contributes to previous parametric classifiers’ degraded performance in recognising ‘New’ classes.

Setting. Consider $f(x)$ as the representation fed to the classifier, *decoupled* training, as the previous settings adopted, indicates $f(x)$ is decoupled when computing the logits l , thus the classification objective won’t supervise representation learning. While for *joint* training, the representations are jointly optimised by classification.

Result & discussion. The results are illustrated in Fig. 4. When adopting the self-labelling strategy, there is a sharp drop in ‘Old’ class performance on both datasets, while for the ‘New’ classes, it can improve by 13 points on CIFAR100, and drop by a small margin on CUB. In contrast,

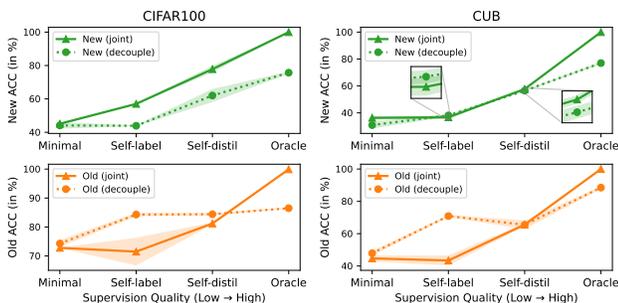


Figure 4. **Results with different training paradigms.** *Decouple* denotes the classifier adopts decoupled features, while *joint* indicates the classification objective can affect representation learning. Joint training is helpful when high-quality supervision is available, otherwise, it could lead to degraded representations.

when a stronger pseudo-labelling strategy (self-distillation) or even oracle labels are utilised, we observe consistent gains from joint training. This means that *the joint training strategy does not necessarily result in UNO [16]’s low performance in ‘New’ classes*; on the contrary, it can even boost ‘New’ class performance by a notable margin. Our overall explanation is that UNO’s framework could not make reliable pseudo-labels, thus restricting its ability to benefit from joint training. The joint training strategy is not to blame and is, in fact, helpful. When switching to a more advanced pseudo-labelling paradigm that produces higher-quality pseudo-labels, the help from joint training can be even more significant.

3.4. The Devil Is in the Biased Predictions

Motivation. In Secs. 3.2 and 3.3, we verified the effectiveness of two design choices when high-quality pseudo labels are available, and concluded the key to previous work’s degraded performance is unreliable pseudo labels. We then further diagnose the statistics of its predictions as follows.

Setting. We categorise the model’s errors into four types: ‘True Old’, ‘False New’, ‘False Old’, and ‘True New’ according to the relationship between predicted and ground-truth class. E.g., ‘True New’ refers to predicting a ‘New’ class sample to another ‘New’ class, while ‘False Old’ indicates predicting a ‘New’ class sample as some ‘Old’ class.

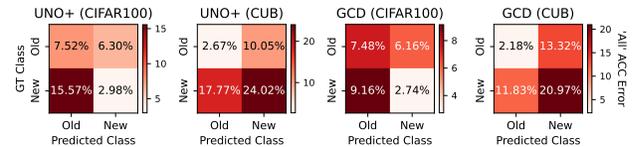


Figure 5. **Prediction bias between ‘Old’/‘New’ classes.** We simplify the results to binary classification and categorise errors in ‘All’ ACC into four types. Both works, especially UNO+, are prone to make ‘False Old’ predictions, and many samples corresponding to ‘New’ classes are misclassified as an ‘Old’ class.

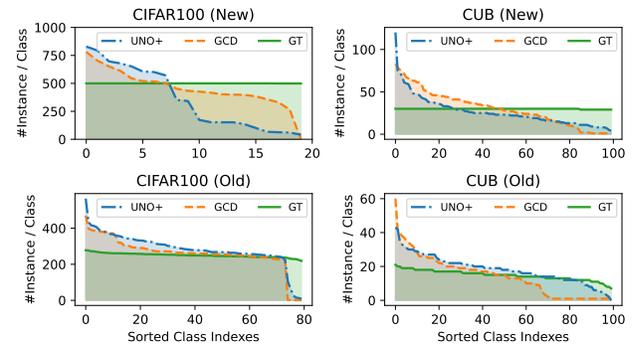


Figure 6. **Prediction bias across ‘Old’/‘New’ classes.** We show the per-class prediction distributions. Both works, especially UNO+, are prone to make biased predictions. Across all classes, the predictions are unexpectedly biased towards the head classes.

Result & discussion. We observe two types of prediction bias. In Fig. 5, both works, especially UNO+ [16], are prone to make “False Old” predictions. In other words, their predictions are biased towards ‘Old’ classes. Besides, the “True New” errors are also notable, indicating that misclassification within ‘New’ classes is also common. We then depict the predictions’ overall distribution across ‘Old’/‘New’ classes in Fig. 6, and both works show highly biased predictions. This double-bias phenomenon then motivated the prediction entropy regularisation design in our method.

4. Method

In this section, we present the whole picture of this simple yet effective method (see Fig. 7), a one-stage framework that builds on GCD [39], and jointly trains a parametric classifier with self-distillation and entropy regularisation. And in Sec. 5.3, we discuss the step-by-step changes that lead a simple baseline to our solution.

4.1. Representation Learning

Our representation learning objective follows GCD [39], which is supervised contrastive learning [24] on labelled samples, and self-supervised contrastive learning [9] on all samples. Formally, given two views (random augmentations) x_i , and x'_i of the same image in a mini-batch B , the self-supervised contrastive loss is written as:

$$\mathcal{L}_{\text{rep}}^u = \frac{1}{|B|} \sum_{i \in B} -\log \frac{\exp(\mathbf{z}_i^\top \mathbf{z}'_i / \tau_u)}{\sum_{i \neq n} \exp(\mathbf{z}_i^\top \mathbf{z}'_n / \tau_u)}, \quad (1)$$

where the feature $\mathbf{z}_i = g(f(x_i))$ and is ℓ_2 -normalised, f, g denote the backbone and the projection head, and τ_u is a temperature value. The supervised contrastive loss is similar, and the major difference is that positive samples are matched by their labels, formally written as:

$$\mathcal{L}_{\text{rep}}^s = \frac{1}{|B^l|} \sum_{i \in B^l} \frac{1}{|\mathcal{N}_i|} \sum_{q \in \mathcal{N}_i} -\log \frac{\exp(\mathbf{z}_i^\top \mathbf{z}'_q / \tau_c)}{\sum_{i \neq n} \exp(\mathbf{z}_i^\top \mathbf{z}'_n / \tau_c)}, \quad (2)$$

where \mathcal{N}_i indexes all other images in the same batch that hold the same label as x_i . The overall representation learning loss is balanced with λ : $\mathcal{L}_{\text{rep}} = (1 - \lambda)\mathcal{L}_{\text{rep}}^u + \lambda\mathcal{L}_{\text{rep}}^s$, where B^l corresponds to the labelled subset of B .

4.2. Parametric Classification

Our parametric classification paradigm follows the self-distillation [8, 2] fashion. Formally, with $K = |\mathcal{Y}_l \cup \mathcal{Y}_u|$ denoting the total number of categories, we randomly initialise a set of prototypes $\mathcal{C} = \{c_1, \dots, c_K\}$, each standing for one category. During training, we calculate the soft label for each augmented view x_i by softmax on cosine similarity between the hidden feature $\mathbf{h}_i = f(x_i)$ and the prototypes

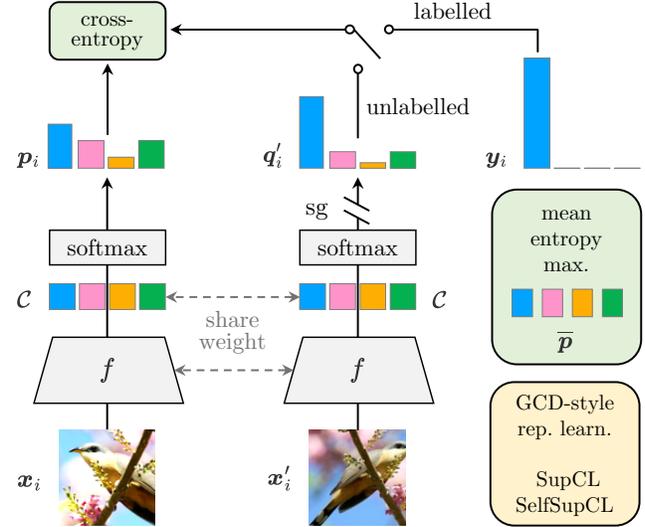


Figure 7. **The overall framework of our method.** For unlabelled samples, the pseudo-labels are from sharpened predictions of another random augmented view. And for labelled samples, we simply adopt the ground truth. Details for representation learning and the mean-entropy-maximisation regulariser are omitted for simplicity, and please refer to the text. (Also see Fig. 1 for a high-level comparison with previous works)

\mathcal{C} scaled by $1/\tau_s$:

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

and the soft pseudo-label q'_i is produced by another view x_i with a sharper temperature τ_t in a similar fashion. The classification objectives are then simply cross-entropy loss $\ell(q', p) = -\sum_k q'^{(k)} \log p^{(k)}$ between the predictions and pseudo-labels or ground-truth labels:

$$\mathcal{L}_{\text{cls}}^u = \frac{1}{|B|} \sum_{i \in B} \ell(q'_i, p_i) - \varepsilon H(\bar{p}), \quad \mathcal{L}_{\text{cls}}^s = \frac{1}{|B^l|} \sum_{i \in B^l} \ell(y_i, p_i), \quad (4)$$

where y_i denote the one-hot label of x_i . We also adopt a mean-entropy maximisation regulariser [2] for the unsupervised objective. Here $\bar{p} = \frac{1}{2|B|} \sum_{i \in B} (p_i + p'_i)$ denotes the mean prediction of a batch, and the entropy $H(\bar{p}) = -\sum_k \bar{p}^{(k)} \log \bar{p}^{(k)}$. Then the classification objective is $\mathcal{L}_{\text{cls}} = (1 - \lambda)\mathcal{L}_{\text{cls}}^u + \lambda\mathcal{L}_{\text{cls}}^s$, and the overall objective is simply $\mathcal{L}_{\text{rep}} + \mathcal{L}_{\text{cls}}$.

Discussions. Please note that this work doesn’t aim to promote new methods but to examine existing solutions, provide insights into their failures and build a simple yet strong baseline solution. The paradigm of producing pseudo-labels from sharpened predictions of another augmented view appears to resemble consistency-based methods [33, 4, 36] in the SSL community. However, despite differences in aug-

mentation strategies and soft/hard pseudo-labels, our approach jointly performs category discovery and self-training style learning, while the SSL methods purely focus on bootstrapping itself with unlabelled data, and does not discover novel categories. Besides, entropy regularisation is also explored in deep clustering to avoid trivial solution [2]. In contrast, our method shows its help in overcoming the prediction bias between and within seen/novel classes (Figs. 9 and 10), and enforcing robustness to unknown numbers of categories (Fig. 11).

5. Experiments

5.1. Experimental Setup

Datasets. We validate the effectiveness of our method on the generic image recognition benchmark (including CIFAR10/100 [26] and ImageNet-100 [37]), the recently proposed Semantic Shift Benchmark [40] (SSB, including CUB [41], Stanford Cars [25], and FGVC-Aircraft [28]), and the harder Herbarium 19 [35] and ImageNet-1K [13]. For each dataset, we follow [39] to sample a subset of all classes as the labelled (‘Old’) classes \mathcal{Y}_l ; 50% of the images from these labelled classes are used to construct \mathcal{D}^l , and the remaining images are regarded as the unlabelled data \mathcal{D}^u . See Tab. 1 for statistics of the datasets we evaluate on.

Dataset	Balance	Labelled		Unlabelled	
		#Image	#Class	#Image	#Class
CIFAR10 [26]	✓	12.5K	5	37.5K	10
CIFAR100 [26]	✓	20.0K	80	30.0K	100
ImageNet-100 [37]	✓	31.9K	50	95.3K	100
CUB [41]	✓	1.5K	100	4.5K	200
Stanford Cars [25]	✓	2.0K	98	6.1K	196
FGVC-Aircraft [28]	✓	1.7K	50	5.0K	50
Herbarium 19 [35]	✗	8.9K	341	25.4K	683
ImageNet-1K [13]	✓	321K	500	960K	1000

Table 1. Statistics of the datasets we evaluate on.

Evaluation protocol. We evaluate the model performance with clustering accuracy (ACC) following standard practice [39]. During evaluation, given the ground truth y^* and the predicted labels \hat{y} , the ACC is calculated as $ACC = \frac{1}{M} \sum_{i=1}^M \mathbb{1}(y_i^* = p(\hat{y}_i))$ where $M = |\mathcal{D}^u|$, and p is the optimal permutation that matches the predicted cluster assignments to the ground truth class labels.

Implementation details. Following GCD [39], we train all methods with a ViT-B/16 backbone [14] pre-trained with DINO [8]. We use the output of [CLS] token with a dimension of 768 as the feature for an image, and only fine-tune the last block of the backbone. We train with a batch size of 128 for 200 epochs with an initial learning rate of 0.1 decayed with a cosine schedule on each dataset. Aligning with [39], the balancing factor λ is set to 0.35, and the temperature values τ_u, τ_c as 0.07, 1.0, respectively. For the classification objective, we set τ_s to 0.1, and τ_t is initialised

to 0.07, then warmed up to 0.04 with a cosine schedule in the starting 30 epochs. All experiments are done with an NVIDIA GeForce RTX 3090 GPU.

5.2. Comparison With the State of the Arts

We compare with state-of-the-art methods in generalized category discovery (ORCA [5] and GCD [39]), strong baselines derived from novel category discovery (RS+ [20] and UNO+ [16]), and k -means [27] on DINO [8] features. On both the fine-grained SSB benchmark (Tab. 2) and generic image recognition datasets (Tab. 3), our method achieves notable improvements in recognising ‘New’ classes (across the instances in \mathcal{D}^u that belong to classes in $\mathcal{Y}_u \setminus \mathcal{Y}_l$), outperforming the SOTAs by around 10%. The results in old classes are also competing against the best-performing baselines. Given that the ability to discover ‘New’ classes is a more desirable ability, the results are quite encouraging.

In Tab. 4, we also report the results on Herbarium 19 [35], a naturally long-tailed fine-grained dataset that is closer to the real-world application of generalized category discovery, and ImageNet-1K [13], a large-scale generic classification dataset. Still, our method shows consistent improvements in all metrics.

Methods	CUB			Stanford Cars			FGVC-Aircraft		
	All	Old	New	All	Old	New	All	Old	New
k -means [27]	34.3	38.9	32.1	12.8	10.6	13.8	16.0	14.4	16.8
RS+ [20]	33.3	51.6	24.2	28.3	61.8	12.1	26.9	36.4	22.2
UNO+ [16]	35.1	49.0	28.1	35.5	70.5	18.6	40.3	56.4	32.2
ORCA [5]	35.3	45.6	30.2	23.5	50.1	10.7	22.0	31.8	17.1
GCD [39]	51.3	56.6	48.7	39.0	57.6	29.9	45.0	41.1	46.9
SimGCD	60.3	65.6	57.7	53.8	71.9	45.0	54.2	59.1	51.8
Δ	+9.0	+9.0	+9.0	+14.8	+14.3	+15.1	+9.2	+18.0	+4.9

Table 2. Results on the Semantic Shift Benchmark [40].

Methods	CIFAR10			CIFAR100			ImageNet-100		
	All	Old	New	All	Old	New	All	Old	New
k -means [27]	83.6	85.7	82.5	52.0	52.2	50.8	72.7	75.5	71.3
RS+ [20]	46.8	19.2	60.5	58.2	77.6	19.3	37.1	61.6	24.8
UNO+ [16]	68.6	98.3	53.8	69.5	80.6	47.2	70.3	95.0	57.9
ORCA [5]	81.8	86.2	79.6	69.0	77.4	52.0	73.5	92.6	63.9
GCD [39]	91.5	97.9	88.2	73.0	76.2	66.5	74.1	89.8	66.3
SimGCD	97.1	95.1	98.1	80.1	81.2	77.8	83.0	93.1	77.9
Δ	+5.6	-2.8	+9.9	+7.1	+5.0	+11.3	+8.9	+3.3	+11.6

Table 3. Results on generic image recognition datasets.

Methods	Herbarium 19			ImageNet-1K		
	All	Old	New	All	Old	New
k -means [27]	13.0	12.2	13.4	-	-	-
RS+ [20]	27.9	55.8	12.8	-	-	-
UNO+ [16]	28.3	53.7	14.7	-	-	-
ORCA [5]	20.9	30.9	15.5	-	-	-
GCD [39]	35.4	51.0	27.0	52.5	72.5	42.2
SimGCD	44.0	58.0	36.4	57.1	77.3	46.9
Δ	+8.6	+7.0	+9.4	+4.6	+4.8	+4.7

Table 4. Results on more challenging datasets.

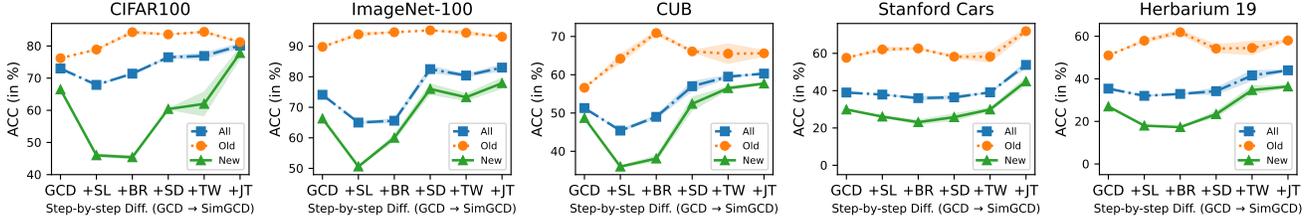


Figure 8. **Step-by-step differences from GCD [39] to SimGCD.** (SL: self-labelling, BR: post-backbone representation, SD: self-distillation, TW: teacher temperature warmup, JT: joint training)

Methods	CF100	CUB	Herb19	IN-100	IN-1K
GCD [39]	7.5m	9m	2.5h	36m	7.7h
SimGCD	1m	18s	3.5m	9.5m	0.6h

Table 5. Inference time over the unlabelled split.

In Tab. 5, we compare the inference time with GCD [39], one iconic non-parametric classification method. Let the number of all samples and unlabelled samples be N and N_u , the number of classes K , feature dimension d , and the number of k -means iterations to be t , the time complexity of GCD is $\mathcal{O}(N^2d + NKdt)$ (including k -means++ initialisation), while our method only requires a nearest-neighbour prototype search for each instance, with time complexity $\mathcal{O}(N_u Kd)$. All methods adopt GPU implementations.

5.3. Ablation Study

In Fig. 8, we ablate the key components that bring the baseline method step-by-step to a new SOTA.

Baseline. We start from GCD [39], a non-parametric classification framework. We keep its representation learning objectives unchanged, and first impose the UNO [16]-style self-labelling classification objectives (+SL) to it, thus transforming it into a parametric classifier. The classifier is built on the projector, and detached from representation learning. Results on ‘Old’ classes generally improve, while results on ‘New’ classes see a sharp drop. This is expected due to UNO’s strong bias toward ‘Old’ classes (Fig. 5).

Improving the representations. As suggested in Sec. 3.2, we build the classifier on the backbone (+BR). This further makes notable improvements in ‘Old’ classes, while changes in ‘New’ classes vary across datasets. This indicates that the pseudo labels’ quality is insufficient to benefit from the post-backbone representations (Fig. 3).

Improving the pseudo labels. We start by replacing the self-labelling strategy with our self-distillation paradigm. As shown in column (+SD), we achieve consistent improvements across all datasets by a large margin (e.g., 26% in CIFAR100, 13% in CUB) in ‘New’ classes. We then further adopt a teacher temperature warmup strategy (+TW) to lower the confidence of the pseudo-labels at an earlier stage. The intuition is that at the beginning, both the classifier and the representation are not well fitted to the target data, thus the pseudo-labels are not quite reliable. This is shown to

be helpful for fine-grained classification datasets, while for generic classification datasets, which are similar to the pre-training data (ImageNet), the unreliable pseudo label is not a problem, thus lowering the confidence does not show help. For simplicity, we keep the training strategy consistent.

Jointly training the representation. Previous settings adopt a decoupled training strategy for consistent representations with GCD [39] and fair comparison. Finally, as confirmed in Sec. 3.3, we jointly supervise the representation with the classification objective (+JT). This results in a consistent improvement in ‘New’ classes for all datasets. Changes in ‘Old’ classes are mostly neutral or positive, with a notable drop in CIFAR100. Our intuition is that the original representations are already good enough for ‘Old’ classes in this dataset, and some incorrect pseudo labels lead to slight degradation in this case.

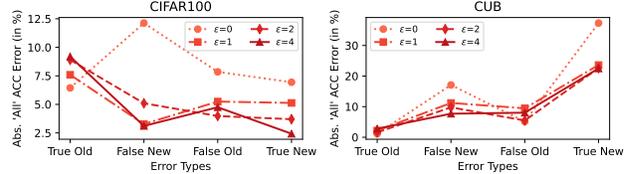


Figure 9. **Effect of entropy regularisation on four types of classification errors.** Appropriate entropy regularisation helps overcome the bias between ‘Old’/‘New’ classes (see ‘False New’ and ‘False Old’, lower is better).

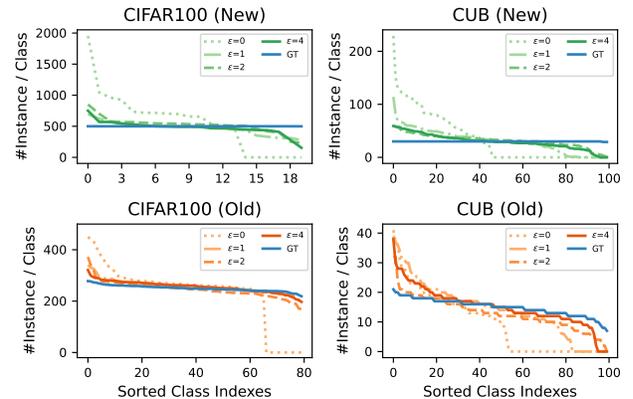


Figure 10. **Per-class prediction distributions with different entropy regularisation weights.** Proper entropy regularisation helps overcome the bias across ‘Old’/‘New’ classes, and approach the GT class distribution.

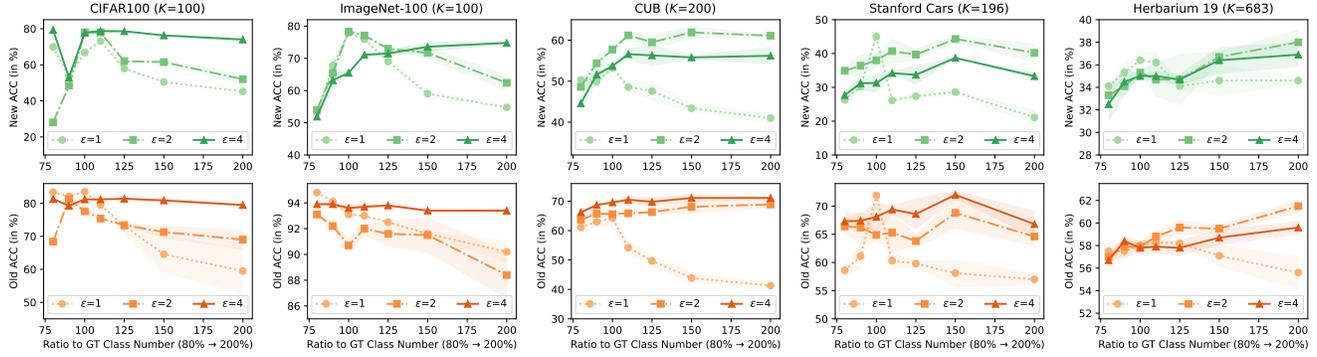


Figure 11. **Results with different numbers of categories.** Stronger entropy regularisation effectively enforces the model’s robustness to unknown numbers of categories, but over-regularisation may limit the ability to recognise ‘New’ classes under ground-truth class numbers.

5.4. Analyses And Discussions

Entropy regularisation helps overcome prediction bias.

We verify the effectiveness of entropy regularisation in overcoming prediction bias by diagnosing the model’s classification errors and class-wise prediction distributions. Fig. 9 shows that this term consistently helps reduce “False New” and “False Old” errors, which refer to predicting an ‘Old’ class sample to a ‘New’ class, and vice-versa. Besides, Fig. 10 shows proper entropy regularisation helps overcome the imbalanced pseudo labels across all classes, and approach the ground truth (GT) class distribution.

Entropy regularisation enforces robustness to unknown numbers of categories. The main text assumed the category number is known *a-priori* following prior works [20, 46, 51, 16], which is impractical [49]. In Fig. 11, we present the results with different numbers of categories on five representative datasets. A category number lower than the ground truth significantly limits the ability to discover ‘New’ categories, and the model tends to focus more on the ‘Old’ classes. On the other hand, increasing the category number results in less harm to the generic image recognition datasets and can even be helpful for some datasets. When a stronger entropy penalty is imposed, the model shows strong robustness to the category number. Interestingly, further analysis in Fig. 12 shows the network prefers to keep the number of active prototypes low and close to the real category number. This finding is inspiring and could ease the deployment of GCD in real-world scenarios.

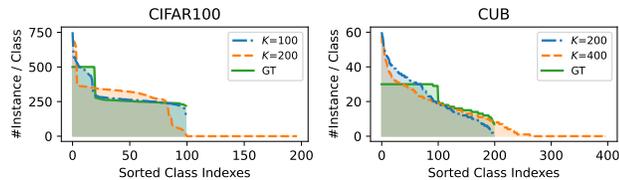


Figure 12. **Per-class prediction distributions with different numbers of categories.** Our method effectively identifies the criterion for ‘New’ classes, thus keeping the number of active prototypes close to the ground-truth class number.

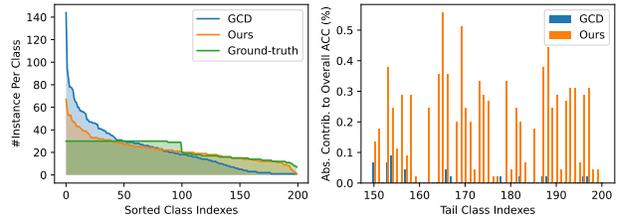


Figure 13. **Prediction analysis against GCD [39].** Left: Based on identical representations, the non-parametric classifier (semi-supervised *k*-means) adopted by GCD produces highly imbalanced predictions, while our method better fits the true distribution; Right: our method significantly improves GCD’s tail classes.

What makes for the significant improvements over GCD given identical representations?

One interesting message from Fig. 8 is that, even with the same representations (col. +TW), we can already improve GCD by a large margin. We thus study the classification predictions and the major components that lead to the performance gap. As shown in Fig. 13, the non-parametric classifier (semi-supervised *k*-means) adopted by GCD [39] produces highly imbalanced predictions, while our method better fits the true distribution. Further analysis (right part) shows that our method significantly improves over the tail classes of GCD.

How does the classification objective change the representations?

In Fig. 8, we have shown that jointly training the representations with the classification objective can lead to $\sim 15\%$ boost in ‘New’ classes on CIFAR100. We study this difference by visualising the representations before and after tuning with t-SNE [38]. As in Fig. 14, jointly tuning the feature leads to less ambiguity, larger margins, and compacter clusters. Concerning why this is not as helpful for CUB: we hypothesise that one important factor lies in how transferable the features learned in ‘Old’ classes are to ‘New’ classes. While it may be easier for a cat classifier to be adapted to dogs, things can be different for fine-grained bird recognition. Besides, the small scale of CUB, which contains only 6k images while holding a large class split (200), might also make it hard to learn transferable features.

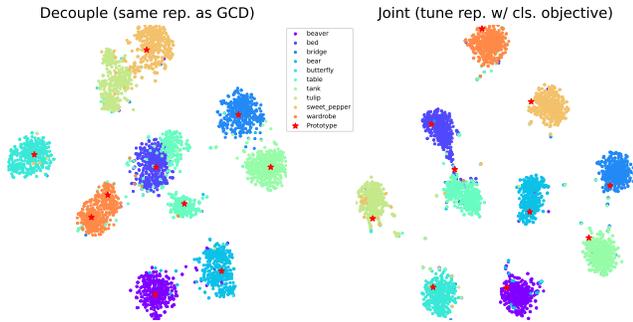


Figure 14. T-SNE [38] visualisation of the representations of 10 classes randomly sampled from CIFAR100 [26]. Jointly supervising representation learning with a classification objective helps disambiguate (e.g., bed & table) and forms compacter clusters.

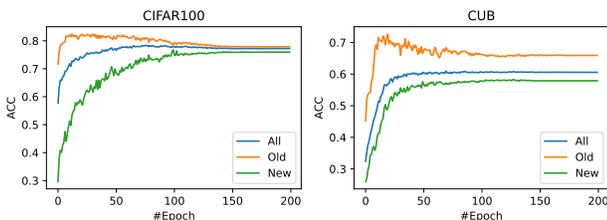


Figure 15. Performance evolution throughout the model learning process. We observe a trade-off between the performance in ‘Old’ and ‘New’ categories, which is common across datasets.

Trade-off between ‘Old’ and ‘New’ categories. We plot the performance evolution throughout the model learning process in Fig. 15. It can be observed that the performance on the ‘Old’ categories first climbs to the highest point at the early stage of training and then slowly degrades as the performance on the ‘New’ categories improves. We believe this demonstrates an important aspect of the design of models for the GCD problem: the performance on the ‘Old’ categories may be in odd with the performance on the ‘New’ categories, how to achieve a better trade-off between these two could be an interesting investigation for future works.

6. Limitations and Potential Future Works

Representation learning. This paper mainly targets improving the classification ability for generalized category discovery. The representation learning, however, follows the prior work GCD [39]. It is expectable that the quality of representation learning can be improved. For instance, generally, by using more advanced geometric and photometric data augmentations [18], and even multiple local crops [7]. Further, can the design of data augmentations be better aligned with the classification criteria of the target data? For another example, using a large batch size has been shown to be critical to the performance of contrastive learning-based frameworks [9]. However, the batch size adopted by GCD [39] is only 128, which might limit the quality of learned representations. Moreover, is the super-

vised contrastive learning plus self-supervised contrastive learning paradigm the ultimate answer to form the feature manifold? We believe that advances in representation learning can lead to further gains.

Alignment to human-defined categories. This paper follows the common practice of previous works where human labels in seen categories implicitly define the metric for unseen ones, which can be viewed as an effort to align algorithm-discovered categories with human-defined ones. However, labels in seen categories may not be good guidance when there is a gap between seen ones and the novel categories we want to discover, e.g., how to use the labelled images in ImageNet to discover novel categories in CUB? For another example, when we use a very big class vocabulary (e.g., the full ImageNet-22K [13]), categories could overlap with each other, and be in different granularities. Further, assigning text names to the discovered categories still requires a matching process, what if further utilising the relationship between class names, and directly predicting the novel categories in the text space? We believe the alignment between algorithm-discovered categories and human-defined categories is of high research value for future works.

Ethical considerations. Current methods commonly suffer from low-data or long-tailed scenarios. Depending on the data and classification criteria of specific tasks, discrimination against minority categories or instances is possible.

7. Conclusion

This study investigates the reasons behind the failure of previous parametric classifiers in recognizing novel classes in GCD and uncovers that unreliable pseudo-labels, which exhibit significant biases, are the crucial factor. We propose a simple yet effective parametric classification method that addresses these issues and achieves state-of-the-art performance on multiple GCD benchmarks. Our findings provide insights into the design of robust classifiers for discovering novel categories and we hope our proposed framework will serve as a strong baseline to facilitate future studies in this field and contribute to the development of more accurate and reliable methods for category discovery.

Acknowledgements

This work has been supported by Hong Kong Research Grant Council - Early Career Scheme (Grant No. 27209621), General Research Fund Scheme (Grant No. 17202422), and RGC Matching Fund Scheme (RMGS). Part of the described research work is conducted in the JC STEM Lab of Robotics for Soft Materials funded by The Hong Kong Jockey Club Charities Trust. The authors acknowledge SmartMore and MEGVII for partial computing support, and Zhisheng Zhong for professional suggestions.

References

- [1] Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR*, 2020. 3
- [2] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Florian Bordes, Pascal Vincent, Armand Joulin, Mike Rabbat, and Nicolas Ballas. Masked siamese networks for label-efficient learning. In *ECCV*, 2022. 2, 3, 5, 6
- [3] Mahmoud Assran, Mathilde Caron, Ishan Misra, Piotr Bojanowski, Armand Joulin, Nicolas Ballas, and Michael Rabbat. Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *ICCV*, 2021. 3
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *NeurIPS*, 2019. 1, 2, 5
- [5] Kaidi Cao, Maria Brbić, and Jure Leskovec. Open-world semi-supervised learning. In *ICLR*, 2022. 1, 2, 6
- [6] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep Clustering for Unsupervised Learning of Visual Features. In *ECCV*, 2018. 3
- [7] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *nips*, 2020. 3, 9
- [8] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2, 3, 5, 6
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 3, 5, 9
- [10] Yanbei Chen, Xiatian Zhu, Wei Li, and Shaogang Gong. Semi-supervised learning under class distribution mismatch. In *AAAI*, 2020. 2
- [11] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *CVPR*, 2021. 3
- [12] Quan Cui, Bingchen Zhao, Zhao-Min Chen, Borui Zhao, Renjie Song, Jiajun Liang, Boyan Zhou, and Osamu Yoshie. Discriminability-transferability trade-off: An information-theoretic perspective. In *ECCV*, 2022. 4
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 6, 9
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 6
- [15] Yixin Fei, Zhongkai Zhao, Siwei Yang, and Bingchen Zhao. Xcon: Learning with experts for fine-grained category discovery. In *BMVC*, 2022. 2
- [16] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *ICCV*, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [17] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations. In *ICLR*, 2018. 2
- [18] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: a new approach to self-supervised learning. In *NeurIPS*, 2020. 9
- [19] Lan-Zhe Guo, Zhen-Yu Zhang, Yuan Jiang, Yu-Feng Li, and Zhi-Hua Zhou. Safe deep semi-supervised learning for unseen-class unlabeled data. In *ICML*, 2020. 2
- [20] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Autonovel: Automatically discovering and learning novel visual categories. *IEEE TPAMI*, 2021. 1, 2, 3, 6, 8
- [21] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *ICCV*, 2019. 2, 3
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1
- [24] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, 2020. 3, 5
- [25] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV Workshops*, 2013. 6
- [26] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical Report*, 2009. 6, 9
- [27] James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1967. 6
- [28] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 6
- [29] Avital Oliver, Augustus Odena, Colin Raffel, Ekin D Cubuk, and Ian J Goodfellow. Realistic evaluation of deep semi-supervised learning algorithms. In *NeurIPS*, 2018. 1, 2
- [30] Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Semi-supervised learning with scarce annotations. In *CVPR Workshops*, 2020. 2
- [31] Subhankar Roy, Mingxuan Liu, Zhun Zhong, Nicu Sebe, and Elisa Ricci. Class-incremental novel class discovery. In *ECCV*, 2022. 3
- [32] Kuniaki Saito, Donghyun Kim, and Kate Saenko. Open-match: Open-set semi-supervised learning with open-set consistency regularization. In *NeurIPS*, 2021. 2

- [33] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *NeurIPS*, 2020. 1, 2, 5
- [34] Yiyu Sun and Yixuan Li. Opencon: Open-world contrastive learning. *TMLR*, 2023. 2
- [35] Kiat Chuan Tan, Yulong Liu, Barbara Ambrose, Melissa Tulig, and Serge Belongie. The herbarium challenge 2019 dataset. *arXiv preprint arXiv:1906.05372*, 2019. 6
- [36] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017. 2, 5
- [37] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *ECCV*, 2020. 6
- [38] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 2008. 8, 9
- [39] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 7, 8, 9
- [40] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: A good closed-set classifier is all you need? In *ICLR*, 2022. 6
- [41] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. Caltech-UCSD Birds 200. *Computation & Neural Systems Technical Report*, 2010. 6
- [42] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: L2 hypersphere embedding for face verification. In *ACM MM*, 2017. 3
- [43] Xin Wen, Bingchen Zhao, Anlin Zheng, Xiangyu Zhang, and Xiaojuan Qi. Self-supervised visual representation learning with semantic grouping. In *NeurIPS*, 2022. 3
- [44] Qing Yu, Daiki Ikami, Go Irie, and Kiyoharu Aizawa. Multi-task curriculum framework for open-set semi-supervised learning. In *ECCV*, 2020. 2
- [45] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *ICCV*, 2019. 2
- [46] Bingchen Zhao and Kai Han. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. In *NeurIPS*, 2021. 2, 3, 8
- [47] Bingchen Zhao and Oisín Mac Aodha. Incremental generalized category discovery. In *ICCV*, 2023. 3
- [48] Bingchen Zhao and Xin Wen. Distilling visual priors from self-supervised learning. In *ECCV Workshops*, 2020. 2
- [49] Bingchen Zhao, Xin Wen, and Kai Han. Learning semi-supervised gaussian mixture models for generalized category discovery. In *ICCV*, 2023. 8
- [50] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *CVPR*, 2021. 3
- [51] Zhun Zhong, Linchao Zhu, Zhiming Luo, Shaozi Li, Yi Yang, and Nicu Sebe. Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In *CVPR*, 2021. 2, 3, 8