

# XCon: Learning with Experts for Fine-grained Category Discovery

Yixin Fei<sup>1</sup>

yixin.feiyx@gmail.com

Zhongkai Zhao<sup>1</sup>

zhongkai.zhaok@gmail.com

Siwei Yang<sup>1,3</sup>

swyang.ac@gmail.com

Bingchen Zhao<sup>2,3</sup>

zhaobc.gm@gmail.com

<sup>1</sup> Tongji University

Shanghai, China

<sup>2</sup> University of Edinburgh,

Edinburgh, UK

<sup>3</sup> LunarAI

---

## Abstract

We address the problem of generalized category discovery (GCD) in this paper, *i.e.* clustering the unlabeled images leveraging the information from a set of seen classes, where the unlabeled images could contain both seen classes and unseen classes. The seen classes can be seen as an implicit criterion of classes, which makes this setting different from unsupervised clustering where the cluster criteria may be ambiguous. We mainly concern the problem of discovering categories within a fine-grained dataset since it is one of the most direct applications of category discovery, *i.e.* helping experts discover novel concepts within an unlabeled dataset using the implicit criterion set forth by the seen classes. State-of-the-art methods for generalized category discovery leverage contrastive learning to learn the representations, but the large inter-class similarity and intra-class variance pose a challenge for the methods because the negative examples may contain irrelevant cues for recognizing a category so the algorithms may converge to a local-minima. We present a novel method called Expert-Contrastive Learning (XCon) to help the model to mine useful information from the images by first partitioning the dataset into sub-datasets using  $k$ -means clustering and then performing contrastive learning on each of the sub-datasets to learn fine-grained discriminative features. Experiments on fine-grained datasets show a clear improved performance over the previous best methods, indicating the effectiveness of our method.

## 1 Introduction

Deep learning models have achieved super-human performance on many computer vision problems where large-scale human annotations are available, such as image recognition [1] and object detection [2]. However, collecting a dataset at scales like ImageNet or COCO is not always possible. Consider the scenario of fine-grained recognition such as bird species recognition or medical image analysis, where the annotations require expert knowledge which could be costly to collect, also it is difficult for the collected annotations to cover all the possible classes because new classes keep growing over time.

The problem of generalized category discovery was recently formalized in [3], where the aim is to discover categories within the unlabeled data by leveraging the information




Figure 1:  $k$ -means results on the DINO features and features after fine-tuning with our method. The images in each row represent a cluster in  $k$ -means. The clusters formed by DINO features are mainly based on the class irrelevant cues, *e.g.*, background and object pose. The features learned by our method could cluster the images based on the correct cues, the object classes.

from a set of labeled data. It is assumed that the labeled data contains similar yet distinct classes from the unlabeled data. The labeled data collected by human experts can be seen as an implicit criterion of classes which can be learned by the model to perform clustering on the unlabeled data. This setting is much harder than semi-supervised learning because generalized category discovery does not assume we know all the classes in the data while in semi-supervised learning the assumption is that the labeled data covers all the classes including ones in unlabeled data.

In this paper, we specifically focus on fine-grained generalized category discovery which is a more difficult and practical problem than generic category discovery since field experts are interested in the fine-grained concepts in real applications, and they often have a labeled dataset representing the existing knowledge, so such a fine-grained generalized category discovery method could help them make sense of the unlabeled set by clustering the unlabeled instance according to the criteria implicitly defined in the labeled data. In fine-grained category discovery, the main challenge is the large inter-class similarity and the intra-class variance, different classes may require the model to learn more discriminative features to be able to distinguish, *e.g.*, two different birds could only differ in the beak. We have observed that an unsupervised representation (*e.g.* DINO) could cluster the data based on class irrelevant cues such as the object pose or the background, see the left part of Fig. 1. Based on this observation, we proposed a simple yet effective method to boost the performance of generalized category discovery on fine-grained data named Expert Contrastive Learning (XCon).

In our proposed XCon method, we partition the data into  $k$  expert sub-datasets by directly performing  $k$ -means clustering on self-supervised representations. These  $k$  sub-datasets can be used as a strong prior for the next learning phase because within each of the  $k$  sub-datasets, class-irrelevant cues will be so similar that the model will be forced to learn more class-relevant features within each sub-dataset. Each of these sub-datasets can be viewed as an expert dataset used to eliminate the negative influence introduced by certain kinds of class-irrelevant cues. Each of the projector used to project features of each corresponding sub-dataset is defined as an expert. To learn a robust representation from these datasets, we directly leverage supervised contrastive learning [14] on the labeled data and unsupervised contrastive learning [2] on all the data.

Our contribution is three-fold:

- We observed that self-supervised representations can group the data based on class irrelevant cues which can be exploited to design further methods.
- We proposed a method that can learn discriminative features for fine-grained category discovery by partitioning the data into  $k$  sub-datasets.
- We validated the effectiveness of our proposed method by setting a new state-of-the-art performance on seven tested generalized category discovery benchmarks.

Our code is available at <https://github.com/YiXXin/XCon>.

## 2 Related Works

### 2.1 Novel Category Discovery

Novel Category Discovery (NCD) aims to discover new object categories by transferring the knowledge learned from a set of relevant but different seen classes. This task was first formalized in DTC [1], with earlier works [2, 3] tackling a similar problem. KCL [4] and MCL [5] utilize the pairwise similarity to transfer the clustering model to cross-task scenarios, which can be used to categorize the unseen classes further. A common three-step learning pipeline is proposed in RankStat [6] where the representation is first learned with self-supervision on all the data and then fine-tuned on the labeled data, the final representation used for discovering novel categories is then further fine-tuned using a pair-wise clustering loss on the unlabeled data. Since then, many works [7, 8, 9, 10, 11, 12] begin to focus on this NCD problem and present promising results. Contrastive learning has been explored under this NCD problem by NCL [13], showing strong performance.

Efforts have also been made in extending this problem to the more challenging fine-grained classification scenario by DualRank [14], which leverages the local object parts information to enhance the representations used for discovering novel categories. Our work also focuses on the challenging fine-grained classification scenario. The key difference with prior works is that we use  $k$ -means grouping on a self-supervised feature to provide informative pairs for contrastive learning instead of using MixUp [15] or local object parts [16]. Our work also builds on a newly proposed setting named Generalized Category Discovery (GCD) [17] where the unlabeled examples can come from both seen and unseen classes, which is a more realistic scenario than NCD.

### 2.2 Contrastive Learning

Contrastive learning has been showing to be effective for learning representations [18, 19] in a self-supervised manner using the instance discrimination pretext [20] as the learning objective. Instance discrimination learns the representation by pushing negative examples away from each other and pulling positive examples closer in the embedding space. As informative examples are important for learning representations with contrastive learning, there are works following this direction trying to create more informative negative or positive pairs using MixUp [21, 22] or special augmentations [23].

Our focus is to learn representations that can be used to discover novel fine-grained categories within the unlabeled dataset, for which a strong representation is needed. By creating informative contrastive pairs by partitioning the dataset into  $k$  sub-datasets using  $k$ -means, examples within each sub-dataset will be similar so that the model will be forced to learn more discriminative features. Compared to previous GCD methods with contrastive learning [24], our method shows clear performance improvements.

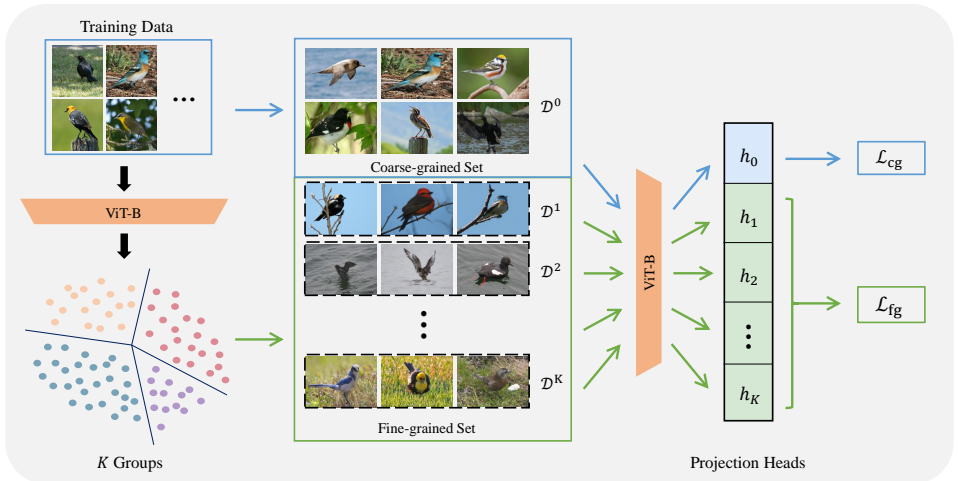


Figure 2: Overview of our XCon framework. We first partition the dataset into  $K$  sub-datasets using  $k$ -means clustering the DINO [10] pretrained representations, then we perform joint contrastive representation learning on each of the partitioned sub-datasets  $\mathcal{D}^1 \dots \mathcal{D}^K$  as well as on the full dataset  $\mathcal{D}^0$ . Each of the partitioned sub-datasets will force the model to learn fine-grained discriminative information, because the background is similar within each of the sub-datasets so the model will need to learn the difference on the objects to be able to distinguish the examples.

### 3 Methods

In GCD, the training dataset contains two parts, a labeled dataset  $\mathcal{D}^l = \{(\mathbf{x}_i^l, y_i^l)\}$  and an unlabeled dataset  $\mathcal{D}^u = \{(\mathbf{x}_i^u, y_i^u)\}$ , where  $y_i^l \in \mathcal{C}^l$  and  $y_i^u \in \mathcal{C}^u$ .  $\mathcal{C}^l$  are only composed of seen classes while  $\mathcal{C}^u$  are composed of both seen and unseen classes, thus  $\mathcal{C}^l \subseteq \mathcal{C}^u$ . The goal of GCD is to learn a model to categorize the instances in  $\mathcal{D}^u$  by leveraging the information from  $\mathcal{D}^l$ . Compared to the previous NCD problem that considers the class sets as  $\mathcal{C}^l \cap \mathcal{C}^u = \emptyset$ , GCD is more challenging and practical.

It has been shown that self-supervised ViT features [10] could be a good initialization for representation learning in GCD [24]. In Vaze *et al.* [24], contrastive learning is used to fine-tune the representation using the information from both labeled and unlabeled datasets, and it is shown that contrastive learning could indeed improve the performance of the representation on the task of GCD. Informative contrastive pairs are important for representation learning, especially in the fine-grained classification setting where the model needs to learn subtle discriminative cues between categories. We proposed a simple method that partitions the dataset into  $k$  sub-datasets by using  $k$ -means on self-supervised features to help construct informative contrastive pairs for representations, the overview of our framework is shown in Fig. 2.

#### 3.1 Preliminary

In this section, we briefly review the method proposed in Vaze *et al.* [24] for GCD, which consists of two parts, representation learning and class assignment. For representation learning, Vaze *et al.* [24] fine-tunes the representation by performing supervised contrastive learning on the labeled data and unsupervised contrastive learning on all the data to avoid overfitting the seen classes.

The unsupervised contrastive loss is defined as

$$\mathcal{L}_i^u = -\log \frac{\exp(\mathbf{z}_i \cdot \hat{\mathbf{z}}_i / \tau)}{\sum_n \mathbb{1}_{[n \neq i]} \exp(\mathbf{z}_i \cdot \mathbf{z}_n / \tau)} \quad (1)$$

where  $\mathbf{z}_i = h(f(x_i))$  is the feature extracted by a backbone  $f(\cdot)$  on the input image  $x_i$  and projected to the embedding space via a projection head  $h(\cdot)$ ,  $\hat{\mathbf{z}}_i$  is the feature from another view of the input image  $\hat{x}_i$ .

The supervised contrastive loss is defined as

$$\mathcal{L}_i^s = -\frac{1}{|\mathcal{N}(i)|} \sum_{q \in \mathcal{N}(i)} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_q / \tau)}{\sum_n \mathbb{1}_{[n \neq i]} \exp(\mathbf{z}_i \cdot \mathbf{z}_n / \tau)} \quad (2)$$

where  $\mathcal{N}(i)$  is the set of indices of images in the minibatch that have the same label  $y_i$  with the anchor image  $i$ .

The final learning objective is the combination of these two losses

$$\mathcal{L}_{\text{coarse}} = (1 - \lambda) \sum_{i \in \mathcal{B}_U \cup \mathcal{B}_L} \mathcal{L}_i^u + \lambda \sum_{i \in \mathcal{B}_L} \mathcal{L}_i^s \quad (3)$$

where  $\lambda$  is used to balance between these two terms,  $\mathcal{B}_U$  is a minibatch of unlabeled images, and  $\mathcal{B}_L$  is a minibatch of labeled images.

For class assignments, the semi-supervised  $k$ -means method is proposed. The overall procedure is similar to the original  $k$ -means method [18], with a key difference that semi-supervised  $k$ -means is aware of the labeled data in  $\mathcal{D}^l$ , and in each step to recompute the cluster assignment, the samples that already have labels will be assigned to the correct cluster regardless of its distance to the nearest cluster centroids.

### 3.2 Dataset Partitioning

The key challenge in representation learning for fine-grained GCD is that the representation is required to be sensitive to the detailed discriminative traits of different classes. Learning the model by contrasting between examples in the full dataset may not help the model to learn such a discriminative representation. Thus, we take advantage of the self-supervised representations that can roughly cluster the images according to the overall image statistics (e.g. background, object pose, etc.) [10] to perform a preprocess on the full dataset by partitioning it into  $k$  expert sub-datasets. The overall statistics within each sub-dataset will be similar and then the model will naturally learn fine-grained discriminative features to distinguish between different examples within each sub-dataset. Each of these expert sub-datasets will be expected to reduce different class-irrelevant cues represented by different overall image statistics.

Specifically, we denote the whole training set as  $\mathcal{D} = \{(x_i, y_i)\}$ . The feature  $\mathbf{v}_i = f(x_i)$  is extracted from each image  $x_i$ . The  $\mathbf{v}_i$  extracted by DINO [10] is incapable of distinguishing between the fine-grained classes since there is no supervision during training, but it will provide a rough description of the image so that similar images will be clustered together. Then, the whole dataset  $\mathcal{D}$  is clustered into  $K$  sub-datasets  $\{\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K\}$  using  $k$ -means, each containing similar images and will be used for fine-grained category discovery later.

### 3.3 Learning discriminative representations

Since the images within each of the partitioned sub-dataset only have fine-grained differences with each other, and each sub-dataset naturally has different global statistics overall, we use a set of projectors  $h_j(\cdot), j = 1, \dots, K$  to project features to each corresponding

sub-spaces in which contrastive learning will be performed. Experts are defined as the  $k$  projectors on the  $k$  sub-datasets which we partition from the whole dataset using  $k$ -means. The expert on each sub-dataset is dedicated to learning fine-grained discriminative features from each sub-dataset. Similar to Vaze *et al.* [24], we apply both supervised contrastive loss and self-supervised contrastive loss to fine-tune the model. Specifically, our proposed fine-grained self-supervised contrastive loss is

$$\mathcal{L}_{\text{fine}}^u = - \sum_{k=1}^K \frac{1}{|\mathcal{B}^k|} \sum_{i \in \mathcal{B}^k} \log \frac{\exp(h_k(\mathbf{v}_i) \cdot h_k(\hat{\mathbf{v}}_i) / \tau)}{\sum_j \mathbb{1}_{[j \neq i]} \exp(h_k(\mathbf{v}_i) \cdot h_k(\mathbf{v}_j) / \tau)} \quad (4)$$

where  $\mathcal{B}^k$  is a minibatch of images sampled from a partitioned dataset  $\mathcal{D}^k$ ,  $\mathbf{v}_i$  and  $\hat{\mathbf{v}}_i$  are two views of one same image through data augmentation, and  $\tau$  is the temperature parameter.

The fine-grained supervised contrastive loss is defined similarly

$$\mathcal{L}_{\text{fine}}^l = - \sum_{k=1}^K \frac{1}{|\mathcal{B}^k|} \sum_{i \in \mathcal{B}^k} \frac{1}{|\mathcal{N}(i)|} \sum_{q \in \mathcal{N}(i)} \log \frac{\exp(h_k(\mathbf{v}_i) \cdot h_k(\mathbf{v}_q) / \tau)}{\sum_j \mathbb{1}_{[j \neq i]} \exp(h_k(\mathbf{v}_i) \cdot h_k(\mathbf{v}_j) / \tau)} \quad (5)$$

where  $\mathcal{N}(i)$  is the set of indices for images with the same label as the anchor image  $i$ .

Thus, the overall loss we propose to learn fine-grained features is the combination of two losses defined above

$$\mathcal{L}_{\text{fine}} = (1 - \lambda) \mathcal{L}_{\text{fine}}^u + \lambda \mathcal{L}_{\text{fine}}^l \quad (6)$$

Together with the loss from Vaze *et al.* [24] defined in Eq. (3), which can be viewed as a coarse-grained loss  $\mathcal{L}_{\text{coarse}}$  compared to our proposed  $\mathcal{L}_{\text{fine}}$ , our optimization objective is

$$\mathcal{L} = \mathcal{L}_{\text{coarse}} + \alpha \mathcal{L}_{\text{fine}} \quad (7)$$

where  $\alpha$  is a parameter to balance between our proposed  $\mathcal{L}_{\text{fine}}$  and the original  $\mathcal{L}_{\text{coarse}}$  from Vaze *et al.* [24]. After the representation is learned, we run the semi-supervised  $k$ -means algorithm to obtain the cluster assignments of each sample.

## 4 Experiments

**Datasets.** We evaluate our method on both generic image classification datasets and fine-grained datasets, with a special focus on the performance of the fine-grained image classification datasets. Following previous works, we choose CIFAR-10/100 [16], ImageNet-100 [8] as the generic image classification datasets. For fine-grained datasets we choose CUB-200 [23], Stanford Cars [13], FGVC-Aircraft [19], and Oxford-IIIT Pet [20]. These fine-grained datasets contain categories from the same *entry level* classes, *e.g.*, birds, cars, aircrafts, and pets. These datasets can be more challenging for GCD methods requiring models to learn highly discriminative features [27]. We split the training data into a labeled dataset and an unlabeled dataset by first dividing all classes equally into a seen class set and an unseen one then sampling 50% images from the seen classes as unlabeled data so that the unlabeled set  $\mathcal{D}^u$  contains images from both seen classes and unseen classes, while the labeled set only contains seen classes, the splits are presented in Table 1.

**Evaluation metric.** We employ the clustering accuracy (ACC) on the unlabeled set to measure the performance. The evaluation metric is defined as below

$$\text{ACC} = \max_{p \in \mathcal{P}(\mathcal{Y}^u)} \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{y_i = p(\hat{y}_i)\} \quad (8)$$

Table 1: Our dataset splits in the experiments.

Dataset		CIFAR10	CIFAR100	ImageNet-100	CUB-200	SCars	Aircraft	Pet
Labelled	Classes	5	80	50	100	98	50	19
	Images	12.5k	20k	31.9k	1498	2000	1666	942
Unlabelled	Classes	10	100	100	200	196	100	37
	Images	37.5k	30k	95.3k	4496	6144	5001	2738

where  $\mathcal{P}$  is the set of all permutations that can match the clustering prediction  $\hat{y}_i$  with the ground-truth label  $y_i$ , we use the Hungarian algorithm [17] to find the best permutation, and  $N$  is the number of images in the unlabeled set. Following [24], we use the metric on three different sets, including ‘All’ referring to the entire unlabeled set  $\mathcal{D}^u$ , ‘Old’ referring to instances in  $\mathcal{D}^u$  belonging to classes in  $\mathcal{C}^l$  and ‘New’ referring to instances in  $\mathcal{D}^u$  belonging to  $\mathcal{C}^u \setminus \mathcal{C}^l$ .

#### 4.1 Implementation details

We follow the implementation of [24] to use ViT-B-16 [5] as the backbone of our method. We initialize the model with the parameters pretrained by DINO [11] on ImageNet and only fine-tune the final transformer block while other blocks are frozen. We implement the projection heads as three layer MLPs following DINO [11], these projection heads will be discarded when testing.

We perform the  $k$ -means to partition the whole dataset into  $k$  sub-datasets. Due to the nature of  $k$ -means clustering, there is an imbalance between different clusters and the number of samples in each sub-dataset is different. The batch size for the entire training dataset is set to 256 and the batch size of all the sub-datasets is set to 32. We create one dataloader for each of eight clusters from  $k$ -means, in each training iteration, 32 images are sampled from each of the eight dataloaders to form a mini-batch with 256 images, so for each mini-batch, the number of images for each clusters is the same.

For the ImageNet dataset, all models are trained for 60 epochs while for other datasets, models are trained for 200 epochs. We set  $\alpha$  to be 0.1 by default. Similar to [24], we use a base learning rate of 0.1 with a cosine annealing schedule and set  $\lambda$  to 0.35. For a fair comparison with existing methods, we use the same semi-supervised  $k$ -means method as [24] to do the evaluation.

#### 4.2 Comparison with the State-of-the-Art

We first compare XCon with the state-of-the-art methods on both generic image classification benchmarks and fine-grained image classification benchmarks. The  $k$ -means method in the tables refers to running  $k$ -means directly on the features extracted from DINO without any further fine-tuning. RankStats+ and UNO+ are two methods modified from two competitive baselines for NCD and adopted to the GCD setting, *i.e.* RankStats [9] and UNO [6].

The results on generic image classification benchmarks are shown in Table 2. On all the datasets we tested, XCon shows the best performance on ‘All’, showing that our method could improve upon previous works. XCon also achieves comparable results with other methods on the other subsets as ‘Old’ and ‘New’. It should be noticed the best performance on ImageNet-100 ‘New’ subset is achieved by naively running a  $k$ -means on DINO features, suggesting that the original features can already represent the unlabeled categories well, and XCon achieves the closest performance compared to this baseline, showing that unlike existing method potentially introducing damage to original feature quality which results in significant performance drop, our method can best preserve the high quality of original features.

Table 2: Results on generic datasets.

Method	CIFAR10			CIFAR100			ImageNet-100		
	All	Old	New	All	Old	New	All	Old	New
$k$ -means [18]	83.6	85.7	82.5	52.0	52.2	50.8	72.7	75.5	<b>71.3</b>
RankStats+	46.8	19.2	60.5	58.2	77.6	19.3	37.1	61.6	24.8
UNO+	68.6	<b>98.3</b>	53.8	69.5	80.6	47.2	70.3	<b>95.0</b>	57.9
GCD [24]	91.5	97.9	88.2	73.0	76.2	<b>66.5</b>	74.1	89.8	66.3
XCon	<b>96.0</b>	97.3	<b>95.4</b>	<b>74.2</b>	<b>81.2</b>	60.3	<b>77.6</b>	93.5	69.7

Table 3: Results on fine-grained datasets.

Method	CUB-200			Stanford-Cars			FGVC-Aircraft			Oxford-Pet		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New
$k$ -means [18]	34.3	38.9	32.1	12.8	10.6	13.8	16.0	14.4	16.8	77.1	70.1	80.7
RankStats+	33.3	51.6	24.2	28.3	61.8	12.1	26.9	36.4	22.2	-	-	-
UNO+	35.1	49.0	28.1	35.5	<b>70.5</b>	18.6	40.3	<b>56.4</b>	32.2	-	-	-
GCD [24]	51.3	<b>56.6</b>	48.7	39.0	57.6	29.9	45.0	41.1	46.9	80.2	85.1	77.6
XCon	<b>52.1</b>	54.3	<b>51.0</b>	<b>40.5</b>	58.8	<b>31.7</b>	<b>47.7</b>	44.4	<b>49.4</b>	<b>86.7</b>	<b>91.5</b>	<b>84.1</b>

We present the results on fine-grained image classification benchmarks in Table 3. Our method shows the best performance on the ‘All’ and ‘New’ with all four datasets we tested while achieving comparable results on ‘Old’. Compared GCD [24] which is the baseline of XCon, our XCon performs better on all metrics on three out of four fine-grained datasets than GCD, and achieves close performance with GCD on the ‘Old’ set on CUB-200, this demonstrates that XCon can indeed improve the performance of the baseline. The UNO+ method outperforms GCD and XCon on ‘Old’ set, the potential reason is pointed out by GCD [24] that a parametric classifier (which UNO+ uses) may be biased towards the classes which have supervision *i.e.* ‘Old’, this is verified by our experiments that UNO+ outperforms XCon on ‘Old’ set but lacks behind on ‘All’ and ‘New’ set.

### 4.3 Ablation study

We perform the ablation study by adjusting each element of our method to inspect the effectiveness of them. For quicker evaluation, we use two fine-grained datasets, *i.e.* CUB-200 and Stanford Cars, and train the model for 100 epochs to ablate the performance.

**Fine-grained and coarse-grained loss.** Table 4 presents the performance of using different combinations of loss terms. We observed that with additional supervision from the coarse-grained loss, the ACC is improved by 3.3-4.0% on CUB-200 and 15.4-28.5% on Stanford Cars. As combining the fine-grained and coarse-grained losses achieves the best performance, it is proved that our proposed method to learn fine-grained features improves GCD methods’ performance in fine-grained benchmarks.

Table 4: Ablation study of fine-grained loss and coarse-grained loss.

$\mathcal{L}_{\text{fine}}$	$\mathcal{L}_{\text{coarse}}$	CUB-200			Stanford-Cars		
		All	Old	New	All	Old	New
✓		48.0	50.5	46.8	21.3	30.6	16.8
	✓	49.9	53.4	48.2	37.1	57.9	27.0
✓	✓	<b>51.8</b>	<b>53.8</b>	<b>50.8</b>	<b>41.0</b>	<b>59.1</b>	<b>32.2</b>



**The weight of fine-grained loss.** We analyze the choice of the weight  $\alpha$  for fine-grained loss in Table 5. We find that XCon can consistently outperform the baseline( $\alpha = 0$ ) with different  $\alpha$ , showing the robust effectiveness of our method. The best result is achieved with  $\alpha = 0.4$  on CUB-200 and with  $\alpha = 0.2$  on Stanford Cars.

Table 5: Ablation study on the weight  $\alpha$  of loss.  $\alpha = 0$  is the baseline(Vaze *et al.* [24]).

$\alpha$	CUB-200			Stanford-Cars		
	All	Old	New	All	Old	New
0	49.9	53.4	48.2	37.1	57.9	27.0
0.1	51.8	53.8	50.8	41.0	59.1	32.2
0.2	51.6	54.5	50.2	<b>42.4</b>	<b>63.0</b>	<b>32.4</b>
0.4	<b>53.4</b>	<b>58.6</b>	<b>50.9</b>	41.1	61.2	31.4

**The number of sub-datasets.** The effect of the sub-dataset number is illustrated in Table 6. Our experiments have shown that *ACC* is increased as long as  $K > 1$ , which shows the usefulness of XCon. Although the performance of XCon is consistently better than the baseline’s, it still varies greatly depending on the number of sub-datasets. When  $K = 2$ , it can reach the highest on the ‘Old’ set, but the lowest on the ‘New’ set, that means with two groups, the overall difference between features is not so great inside each group that the model tends to focus more on the existing coarse-grained knowledge learned from the seen classes. In practice, one could determine the ideal  $k$  by evaluating against a validation set.

Table 6: Ablation study on the number  $K$  of split sub-groups.

$K$	CUB-200			Stanford-Cars		
	All	Old	New	All	Old	New
1	49.9	53.4	48.2	37.1	57.9	27.0
2	51.4	<b>59.3</b>	47.4	40.9	<b>61.0</b>	31.1
4	51.7	54.6	50.2	39.8	55.3	32.3
6	50.3	51.9	49.5	<b>42.1</b>	60.7	<b>33.1</b>
8	<b>51.8</b>	53.8	<b>50.8</b>	41.0	59.1	32.2

We further visualize the feature spaces with TSNE [23] on CIFAR10 by mapping the features into two dimensions for a more direct qualitative analysis. In Fig. 3, we cluster the unlabeled data and compare results of the initial model(DINO) before fine-tuning with ones of our model after fine-tuning of DINO. It is clear that the effect before and after the fine-tuning is significant. DINO can cluster the features into 10 groups roughly, but many samples appear in groups that correspond to other classes. In contrast to DINO, with our model, we can see clear boundaries between different groups, and each group is corresponding to one certain category in CIFAR10.

## 5 Conclusion

In this paper, we propose XCon to address the problem of generalized category discovery with fine-grained image classification benchmarks. XCon first partitions the dataset into  $K$  sub-datasets using  $k$ -means clustering on a self-supervised representation. Each partitioned sub-dataset can be seen as a subset of images that are visually similar and have close coarse-grained representation so that contrastive learning within each of these sub-datasets will force the model to learn fine-grained discriminative features that can help discover fine-

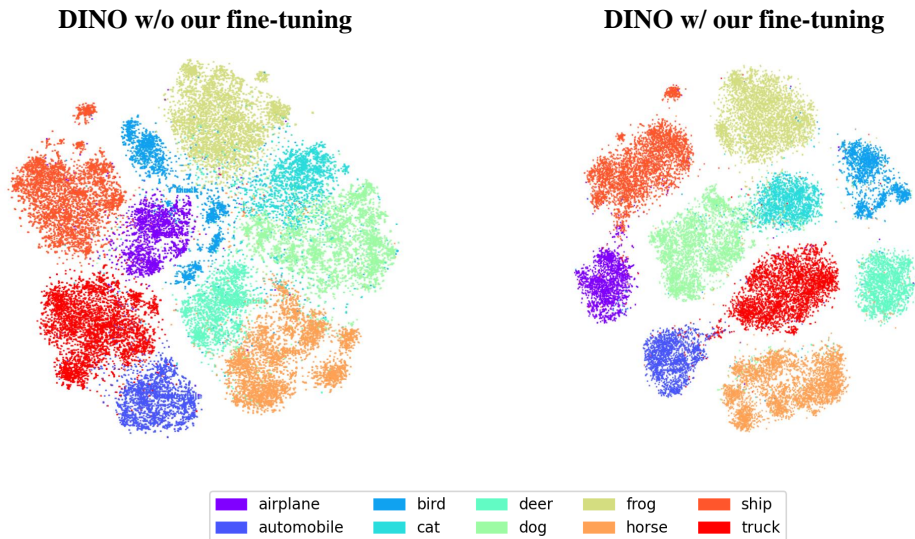


Figure 3: Feature visualization on CIFAR10 with TSNE.

grained categories. Experiments on four fine-grained image classification benchmarks show clear performance improvements of XCon, validating the effectiveness of our method.

## Acknowledge

The authors would like to acknowledge compute support from LunarAI.

## References

- [1] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021.
- [2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [3] Haoang Chi, Feng Liu, Wenjing Yang, Long Lan, Tongliang Liu, Bo Han, Gang Niu, Mingyuan Zhou, and Masashi Sugiyama. Meta discovery: Learning to discover novel classes given very limited data. In *International Conference on Learning Representations*, 2021.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

- [6] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9284–9292, 2021.
- [7] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8401–8409, 2019.
- [8] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *International Conference on Learning Representations*, 2020. URL [https://openreview.net/forum?id=BJl2\\_nVFPB](https://openreview.net/forum?id=BJl2_nVFPB).
- [9] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Autonovel: Automatically discovering and learning novel visual categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [10] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for 393 unsupervised visual representation learning. In *Conference on Computer Vision and Pattern*, volume 394, 2019.
- [11] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. Learning to cluster in order to transfer across domains and tasks. *arXiv preprint arXiv:1711.10125*, 2017.
- [12] Yen-Chang Hsu, Zhaoyang Lv, Joel Schlosser, Phillip Odom, and Zsolt Kira. Multi-class classification without multi-class labels. *arXiv preprint arXiv:1901.00544*, 2019.
- [13] Yannis Kalantidis, Mert Bulent Sariyildiz, Noe Pion, Philippe Weinzaepfel, and Diane Larlus. Hard negative mixing for contrastive learning. *Advances in Neural Information Processing Systems*, 33:21798–21809, 2020.
- [14] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020.
- [15] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013.
- [16] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical report*, 2009.
- [17] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [18] James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1967.
- [19] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.

- [20] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012.
- [21] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [22] Zhiqiang Shen, Zechun Liu, Zhuang Liu, Marios Savvides, Trevor Darrell, and Eric Xing. Un-mix: Rethinking image mixtures for unsupervised visual representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2022.
- [23] Laurens Van Der Maaten. Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research*, 15(1):3221–3245, 2014.
- [24] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7492–7501, 2022.
- [25] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. *Technical Report CNS-TR-2011-001*, California Institute of Technology, 2011.
- [26] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018.
- [27] Bingchen Zhao and Kai Han. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. *Advances in Neural Information Processing Systems*, 34:22982–22994, 2021.
- [28] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10867–10875, 2021.
- [29] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9462–9470, 2021.
- [30] Rui Zhu, Bingchen Zhao, Jingen Liu, Zhenglong Sun, and Chang Wen Chen. Improving contrastive learning by visualizing feature transformation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10306–10315, 2021.