# Adaptive Part Learning for Fine-Grained Generalized Category Discovery: A Plug-and-Play Enhancement

Qiyuan Dai Hanzhuo Huang Yu Wu Sibei Yang<sup>†</sup> School of Information Science and Technology, ShanghaiTech University Wuhan University Sun Yat-sen University

## Abstract

Generalized Category Discovery (GCD) aims to recognize unlabeled images from known and novel classes by distinguishing novel classes from known ones, while also transferring knowledge from another set of labeled images with known classes. Existing GCD methods rely on selfsupervised vision transformers such as DINO for representation learning. However, focusing solely on the global representation of the DINO CLS token introduces an inherent trade-off between discriminability and generalization. In this paper, we introduce an adaptive part discovery and learning method, called APL, which generates consistent object parts and their correspondences across different similar images using a set of shared learnable part queries and DINO part priors, without requiring any additional annotations. More importantly, we propose a novel all-min contrastive loss to learn discriminative yet generalizable part representation, which adaptively highlights discriminative object parts to distinguish similar categories for enhanced discriminability while simultaneously sharing other parts to facilitate knowledge transfer for improved generalization. Our APL can easily be incorporated into different GCD frameworks by replacing their CLS token feature with our part representations, showing significant enhancements on fine-grained datasets.

# 1. Introduction

Trained on large-scale annotated data [9, 25] and simplified by the "closed-world" assumption, where the test data to be classified adheres to predefined training classes, deep learning models [10, 17, 34] have made remarkable progress in image recognition, even surpassing human abilities. However, in practical real-world scenarios, this assumption often fails to hold true. This is not only due to the high expense of fully annotating data but, more importantly, due to the impossibility of predefining all categories because of the inherent variability and openness of the real world, where novel classes emerge alongside the known ones. Therefore, Generalized Category Discovery (GCD) [40] seeks to challenge the "closed-world" assumption and adapt to an "openworld" scenario. In GCD, only a subset of data is labeled with known categories, aiming to assign a category label to each remaining data point, potentially encompassing both known and novel categories.

Two primary challenges arise in GCD: 1) Discriminability: accurately discerning between known and novel classes, and correctly categorizing each; 2) Generalization: transferring shared knowledge from known to novel classes to recognizing novel classes without annotations. Stateof-the-art GCD methods [40, 42, 43] primarily exploit the benefits of generalization representation via self-supervised vision transformers (ViT) like DINO [4] to uphold fundamental generalization capabilities. On one hand, to improve semantic discriminability, some approaches utilize an explicit parametric classifier on known data [43] or learn prompt-adapted backbones [42, 45]. On the other hand, to safeguard generalization capabilities, they introduce entropy regularization to mitigate bias towards tending to predict known categories [43] or incorporate shared dynamic concepts [30], hierarchies [31], or memory bank [45] to enhance shared semantics between known and novel classes.

However, an inherent trade-off arises between discriminability and generalization when utilizing the image semantics at the global level, namely, the feature representation of ViT's [cls] token, in current GCD methods, particularly in open-world fine-grained recognition. *The strong discriminability at the global level impedes the potential for generalization.* The global semantics are rooted in finegrained local semantics, with the [cls] token's representation obtained through self-attention over local patches [4, 10]. Introducing significant differences in the global semantics of similar categories, such as crested auklet and parakeet auklet as shown in Figure 1, can improve discriminability but also lead to similarly notable distinctions in their local semantics (refer to Figure 1b). Consequently, this impedes the sharing and knowledge transfer of local

<sup>&</sup>lt;sup>†</sup>Corresponding author



Figure 1. Comparison of discriminability and generalization between our and existing GCD methods that use global CLS token representation. We (a) allow similar parts in similar classes to be shared to achieve knowledge transfer and generalization, while (d) exploiting discriminative parts to distinguish novel classes from similar known classes. In contrast, the existing GCD methods (b) cannot distinguish different parts within the image, or (c) are easy to classify novel images into similar known classes when using superclass-based generalization. (e) demonstrates the integration of our APL into existing GCD frameworks, and (f) presents the improvements over diverse baselines.

semantics, such as the legs, which should inherently be similar across the two classes. As a contrast, Figure 1a showcases our discriminable yet shareable local semantics. On the contrary, strong generalization on similar categories could impact discriminability on novel categories. Hierarchical or superclass-based generalization may result in similar features among different categories within the same superclass, posing challenges in distinguishing novel categories from their similar known counterparts, as shown in Figure 1c. This occurs because supervised learning on labeled data only emphasizes differences between similar known classes, leaving the discriminative fine-grained cues of novel categories easily overlooked due to the lack of explicit supervision. However, our model can distinguish between similar novel and known classes, such as rhinoceros auklet and crested auklet, thanks to our discriminative local semantics (such as the head), as shown in Figure 1d.

In this study, we aim to improve discriminability while preserving generalization through a unified perspective, namely, adaptive part representation learning (APL): highlighting discriminative fine-grained local semantics to distinguish similar but distinct categories (e.g., created auklet and parakeet auklet distinguishable by their heads), while also allowing for the sharing of their other local semantics (e.g., legs), thereby facilitating knowledge transfer and generalization. Specifically, our APL comprises two main components: 1) Consistent part discovery and correspon*dence*, which learns to decompose images into meaningful semantic parts that can be shared and corresponded one-toone between similar categories, as illustrated in Figure 3c. However, unsupervised object part discovery and correspondence, without part or even object mask annotations, presents challenges. Unfortunately, even though DINO's patch features are capable of perceiving object parts [29], naive image partition yields fragmented regions, lacking

meaningful parts. Similarly, clustering image patches into a fixed number of groups overlooks the complexity of different scenes, including scenes with missing parts due to occlusion and augmentations, resulting in the noisy part discovery and subsequent incorrect correspondences, as shown in Figure 3a. Therefore, unlike prior methods, we identify object parts using a set of learnable part queries, guided by DINO's part prior. Each image's patches are then hardassigned to these queries, forming object parts associated with the same query, as shown in Figure 2. Oueries with empty tokens handles potential missing parts. Additionally, inducing shared part queries from all images enhances robustness compared to single-image clustering and naturally establishes correspondences between parts of different images by using queries as bridges. 2) Diverse and adaptive part representation learning acquires more discriminative semantic parts of fine-grained objects while preserving shared parts for improved generalizability across various classes. To achieve this, we propose two learning objectives: adaptive part contrastive learning and query diversity. Firstly, we propose *all-min* contrastive learning on semantic parts, ensuring that all corresponding parts of positive samples are similar while only constraining the *least* similar parts of negative samples to be dissimilar. This enables the adaptive discovery of discriminative parts, while leaving the generalizable parts of negative samples unconstrained, thereby making object part features shareable and generalizable to similar classes. Next, we constrain the diversity between queries, ensuring that each query can focus on and highlight specific semantic parts and details, thus avoiding collapsing into a global semantic representation.

Our APL which improves GCD through discriminative yet generalizable part discovery without relying on a specific GCD network architecture(refer to Figure 1e), is integrated into the widely used SimGCD [43] as well as the latest SPTNet [42] and CMS [7] to demonstrate its effectiveness. In summary, our contributions are multi-faceted: (1) To the best of our knowledge, we are the first to introduce and differentiate between different object parts for GCD by emphasizing discriminative parts to enhance discriminability, while simultaneously sharing other parts to facilitate knowledge transfer for improved generalization. (2) We propose a simple yet effective unsupervised object part discovery method that seamlessly integrates with GCD frameworks to achieve consistent and corresponding part discovery and segmentation across similar images with the same or different categories. (3) We propose a novel all-min contrastive learning objective to adaptively select and constrain only the most discriminative parts of negative pairs, fostering discriminative yet generalizable representation learning. (4) Our APL is incorporated into the classic SimGCD, recent SPTNet and CMS GCD frameworks, achieving consistent performance improvements across all fine-grained datasets(refer to Figure 1f) and enhancing the average accuracy by 3.7%, 2.9% and 3.7% respectively.

## 2. Related Works

Generalized Category Discovery (GCD) [40] extends the novel category discovery (NCD) to a more realistic setting, where unlabeled data contains samples of both known and novel classes. The pioneering work [40] utilizes the selfsupervised pre-trained DINO [4] as the initialization of the feature extractor, in conjunction with semi-supervised kmeans clustering. It integrates supervised contrastive learning [14, 19, 37] on labeled data for distinguishing known classes, along with unsupervised contrastive learning on different views of both labeled and unlabeled data. Later, DCCL [30] introduces dynamic conception generation to consider inter-class relationships, using infomap to dynamically generate pseudo-labels. SimGCD [43], unlike prior clustering-based approaches, directly employs a parameterized classification head. This circumvents the constraint of clustering methods, which cannot classify individual samples directly, resulting in notable efficiency enhancements. Building upon SimGCD, SPTNet [42] introduces spatial prompt tuning to enable the model to leverage information from local regions, while CMS [7] modifies the classification head for hierarchical clustering and enhances feature generalization through the mean shift technique. Some other methods [44, 46] leverage local-level information, but [44] exploits the structural characteristics of 3D point clouds to facilitate part discovery, whereas [46] fails to obtain semantically meaningful part information.

**Unsupervised Object Part Discovery.** To enhance finegrained image understanding, certain studies advance from the image-level to the object-level and subsequently to an even finer part-level comprehension. Early works [5, 49] such as ProtoPNet [5] explore the utilization of part-level

information to enhance classification. It introduces prototypes into convolutional networks to represent object parts. However, the lack of ground truth or priors for accurate object parts limits the robustness of the part information obtained by these methods [1, 20, 21, 33, 38]. Other methods [24, 28, 48] rely on fully annotated object parts to learn part segmentation, aiming to segment images into distinct regions corresponding to different parts of objects. Recently, with the emergence of DINO [4], some methods use DINO for part discovery. For instance, [8] employ frozen DINO to extract part features for help learning part segmentation. Similarly, VLPart [35] utilizes DINO to establish correspondence between novel parts and base parts, but it still relies on annotations of object parts in base categories. In this paper, by leveraging DINO's coarse part conception, our proposed shared part queries, and novel all-min contrastive loss, we automatically, unsupervised, discover consistent object parts and establish correspondences between them across different images.

# 3. Method

## 3.1. Preliminary

**Problem Statement.** The Generalized Category Discovery (GCD) entails categorizing images during inference, with these images potentially belonging to either labeled known classes or unlabeled novel classes during training. Formally, learning from a labeled set  $\mathcal{D}_l = \{(I_i, c_i)\}_{i=1}^{N_l} \subset \mathcal{I}_l \times C_l$  and an unlabled set  $\mathcal{D}_u = \{I_i\}_{i=1}^{N_u} \subset \mathcal{I}_u$ , each image  $I_i \in \mathbb{R}^{H \times W \times 3}$  in both sets is with height H and width W, the GCD aims to accurately categorizes each image  $I_i \in \mathcal{I}_u$  within the labeled label space  $\mathcal{C}_l$  is a subset of the unlabeled label space  $\mathcal{C}_u$ . Given that the labeled label space  $\mathcal{C}_l$  is a subset of the unlabeled label space  $\mathcal{C}_u$ , GCD not only requires discriminability to distinguish between labeled known classes in  $\mathcal{C}_l$  and remaining novel classes in  $\mathcal{C}_u - \mathcal{C}_l$ , but also transfer shared knowledge from known to novel classes.

DINO Architecture. DINO [4, 29] adopts the standard Vision Transformer (ViT) [11] architecture. Given an image I, DINO initially divides it into a sequence of non-overlapping patch tokens. Each token is then projected into a hidden representation vector via a linear or convolutional layer, resulting in token vectors denoted as  $X_{patch} \in \mathbb{R}^{N \times C}$ , where N is the number of patch tokens, and C specifies vector dimension. To represent the global image feature, an additional learnable CLS token is concatenated to  $X_{patch}$ . Subsequently, all tokens, patch tokens  $X_{\text{patch}}$  and CLS token, integrated with positional embeddings [11], are fed into the Transformer layers that contain several multi-head selfattention [39] (MSA) and feed-forward networks (FFN). Within the multiple transformer layers, the CLS token interacts with image patch tokens  $X_{patch}$  through MSA, aggregating global features. Finally, DINO develops robust



Figure 2. Overall framework of APL. APL achieves consistent and corresponding part discovery, serving as a plug-and-play enhancement.

visual representations, denoted by  $F = [F_{patch}, F_{CLS}]$  with token features  $F_{patch} \in \mathbb{R}^{N \times C}$  and CLS feature  $F_{CLS} \in \mathbb{R}^{1 \times C}$ , achieved through view-invariant self-supervised pretraining [4, 29]. The DINO architecture is formulated as:

$$F_0 = \phi(\text{PE}, [X_{\text{CLS}}, X_{\text{patch}}]),$$

$$[F_{\text{patch}}, F_{\text{CLS}}] \leftarrow F_{l=L}, F_l = \text{FFN}(\text{MSA}(F_{l-1})),$$
(1)

where  $[\cdot, \cdot]$  denotes concatenation, PE represents position encoding,  $\phi$  is the fusion function of position encoding and tokens [4], and l and L represent the index of the l-th transformer layer and the total number of transformer layers, respectively.

**Part Perception in DINO.** With self-supervised pretraining and ViT architecture, DINO exhibits many functionalities that are absent in fully supervised training [11] or ConvNet [18] networks, one of which is the perception of object parts. Previous works [4, 8] find that in the MSA interactions between the CLS token and image patch tokens, the overall attention scores from all M attention heads can roughly reveal the foreground object, while attention scores of each individual head can highlight semantically meaningful parts of the object. Specifically, we use the attention scores from the final-layer MSA of DINO to obtain part attention. In the *m*-th attention head, with the features of the CLS token denoted as  $F_{\text{CLS}}^m \in \mathbb{R}^{1 \times \frac{C}{M}}$  and the features of the patch tokens as  $F_{\text{patch}}^m \in \mathbb{R}^{N \times \frac{C}{M}}$ , the calculation of *m*-th part attention  $A_m$  is as follows:

$$W_m = \text{Softmax}\left(\frac{\text{Proj}_Q^m(F_{\text{CLS}}^m) * \text{Proj}_K^m([F_{\text{CLS}}^m, F_{\text{patch}}^m])^\top}{\sqrt{C}}\right),$$
$$A_m \leftarrow W_m[1, 2: N+1],$$
(2)

where  $\operatorname{Proj}_{Q}^{m}$  and  $\operatorname{Proj}_{K}^{m}$  denote projection matrix of *m*-th head,  $\top$  represents matrix transpose. The  $W_m \in \mathbb{R}^{1 \times (N+1)}$  denotes the attention scores between the CLS token and all tokens, encompassing both the CLS token and patch tokens. Specifically, the attention scores  $A_m \in \mathbb{R}^{1 \times N}$  for patch tokens highlight potential patch regions where parts could be located, essentially constituting the part attention in the *m*-th head.

#### **3.2.** Consistent Part Discovery and Correspondence

The Motivation for Introducing Object Parts is to highlight discriminative object parts for enhanced discriminability, while simultaneously sharing other parts to facilitate knowledge transfer for improved generalization. Since GCD requires learning from unlabeled data, DINO-a selfsupervised, pre-trained model that learns from unlabeled images and their augmentations-is a natural choice for existing GCD methods [42, 43] seeking to obtain generalizable visual representations. Existing GCD methods predominantly use the global semantics of images, particularly the features of DINO's CLS token, for classification or clustering to categorize images. However, these global features may lead to a trade-off between discriminating similar categories and generalizing from similar known to novel ones, as detailed in Sec 1. Therefore, we advocate for integrating local semantics, such as object parts, into GCD. Notably, object parts are not merely naive supplements of finegrained local semantics to global features. More crucially, by delineating the roles of different object parts, discriminative parts can effectively differentiate similar categories, while also allowing for the sharing of other parts across these categories, thereby improving both discriminability and generalization.

Three Principles of Our Part Discovery. (1) Consistency: object parts in images of the same category should be consistent, especially when some images may have missing parts. For instance, in cars, parts like windows, tires, and hoods are common, but due to image augmentation or occlusion, some images may only feature partial parts. Simple single-image-based grouping or clustering on DINO's patch features is highly unreliable, potentially leading to meaningless or inconsistent discoveries of object parts across images, as shown in Figure 3a. (2) Correspondence: some object parts of similar categories should correspond to each other, facilitating the identification of both discriminative and shareable parts. For instance, dogs and cats share common parts like the head, tail, and body. However, only by matching their corresponding parts can we identify the head as potentially more discriminative, while body features may be better suited for sharing and generalizing to other canid species. (3) Efficiency: part discovery should be seamlessly integrated into existing GCD frame-



Figure 3. Part discovery by (a) K-means, (b) DINO priors, and (c) ours. Only ours demonstrates consistency in handling partial parts and correspondences between similar images.

works without the need for additional annotations, external models, or significant computational complexity.

Consistent Part Discovery and Correspondence. To adhere to our third principle of efficiency, we propose directly leveraging DINO, drawing inspiration from its perceptual abilities for object parts (see Sec 3.1), and integrating part discovery into the online learning process of GCD. However, DINO's part perception is too sensitive to variations among different images, thus failing to meet our principles of consistency and correspondence, as shown in Figure 3b. Therefore, instead of solely depending on DINO's part perception, we leverage it as a prior to facilitate learning a set of shared learnable part queries across all images for object discovery, as shown in Figure 2. These cross-image, unified part queries not only ensure more robust object discovery compared to single-image-based methods but also naturally serve as bridges for building correspondence between different images. Specifically, let us denote the set of learnable part queries as  $Q \in \mathbb{R}^{T \times C}$ , where T is the number of queries and C is the dimension number. First, for an input image I, we extract its DINO part priors  $F_{\text{prior}} \in \mathbb{R}^{M \times C}$ based on DINO's M part attentions  $[A_1, A_2, ..., A_M]$  and DINO patch features  $F_{patch}$ , defined in Sec 3.1. Then, we transform the shared learnable queries Q into imagespecific ones  $Q_{I} \in \mathbb{R}^{T \times C}$  via cross-attention between Qand the prior  $F_{prior}$ , as formulated below:

$$F_{\text{prior}} = [\epsilon(A_1) * F_{\text{patch}}, \cdots, \epsilon(A_M) * F_{\text{patch}}],$$
  
$$Q_I = \operatorname{Proj}_{\mathbf{Q}}(Q) * \operatorname{Proj}_{\mathbf{K}}(F_{\text{part}})^\top / \sqrt{C},$$
(3)

where  $\epsilon$  represents a threshold filter function to transform  $A_m$  into a binary matrix to highlight the part regions,  $\operatorname{Proj}_{Q}$  and  $\operatorname{Proj}_{K}$  are the learnable projection matrices. The  $\epsilon(A_m) * F_{\text{patch}}$  represents the feature of the *m*-th part in DINO, and the features from all heads form DINO's part priors  $F_{\text{prior}}$ . Part queries  $Q_{I}$  encode the features of potential object parts in image *I*. Next, object parts are identified

by hard-assigning image patches to the queries based on the similarity of their features  $F_{patch}$  and  $Q_I$ , where patches within the same query constitute a potential object part. As the hard assignment, which is the arg max over similarities, is not differentiable, we approximate the hard assignment and preserve differentiability by Gumbel-Softmax distribution [13, 26], as formulated as follows:

$$H_{g} = \operatorname{softmax}((\frac{\operatorname{Proj}_{Q}(Q_{I}) * \operatorname{Proj}_{K}(F_{\operatorname{patch}})^{\top}}{\sqrt{C}} + G)/\tau),$$
  

$$H_{\operatorname{part}} = (\operatorname{one\_hot}(\arg\max_{N}(H_{g})))^{\top} - \operatorname{sg}(H_{g}) + H_{g},$$
  

$$P_{\operatorname{part}} = H_{\operatorname{part}} * F_{\operatorname{patch}},$$
  
(4)

where each Gumbel noise  $g_n \in G$  with  $g_n$ =  $-\log(-\log(u_n))$  and  $u_n$  are i.i.d. samples drawn from Uniform(0,1),  $\tau$  is the temperature. Proj<sub>o</sub> and  $Proj_{K}$  denote the learnable projection matrices of  $Q_{I}$ and  $F_{\text{patch}}$ . The  $H_{g}$  approximates sampling the one\_hot(arg max<sub>N</sub>(Proj<sub>O</sub>(Q<sub>I</sub>)\*Proj<sub>K</sub>(F<sub>patch</sub>)<sup> $\top$ </sup>/ $\sqrt{C}$ +G)), where  $\arg \max_N$  selects the corresponding patch token with the highest attention score for each part query, and one-hot transforms the patch token index into a one-hot vector. sq represents the stop gradient operator. The resulting binary  $H_{\text{part}} \in \mathbb{R}^{T \times N}$  indicates the hard assignment from N patches to T part queries. Each  $H_t \in \mathbb{R}^{1 \times N}$  within  $H_{\text{part}}$  represents a potential part, with indices containing a value of 1 indicating which patches belong to that specific part. The  $P_t \in \mathbb{R}^{1 \times C}$  within  $P_{\text{part}}$  represents the feature of each part  $H_t$ .

Our query hard assignment for part discovery can adaptively handle consistency and situations where some parts are absent. In such cases, when an object part is too small and its constituent patches are scattered, the corresponding query becomes a meaningless placeholder for the potential missing part. Additionally, it ensures part correspondence across images. For images  $I_a$  and  $I_b$ , their respective parts  $H^a_{\text{part}}$  and  $H^b_{\text{part}}$  are obtained through shared part queries Q. Consequently, parts corresponding to same query, such as  $H^a_t$  and  $H^b_t$  of the t-th query, align with each other.

## 3.3. Diverse and Adaptive Part Learning

While object parts are meant to represent distinct components, without constraints on the learning of randomly initialized learnable parts, they easily correspond to arbitrary image regions rather than meaningful parts or collapse into identical global patches. Additionally, automatically distinguishing discriminative and transferable object parts enables learning their respective representations, crucial for enhancing GCD from both discriminative and generalizable perspectives. Therefore, we propose a novel all-min contrastive learning objective combined with a diverse loss to facilitate part learning.

All-min Contrastive Learning. Its core idea involves (1) constraining all corresponding object parts in a positive pair

to be more similar than non-corresponding pairs, where positive pairs originate from images of the same class in the labeled set or from images and their augmentations in both the labeled and unlabeled sets. This fosters the learning of distinct parts within a class and consistency among parts across images of the same class. (2) In negative pairs of images from different classes, we only enforce dissimilarity on the most dissimilar corresponding part, leaving the similarity of their other corresponding parts unconstrained. This approach capitalizes on the least similar corresponding part's discriminative power in distinguishing between the two images, while permitting other parts to be similar to facilitate the sharing of similar components for generalization. Specifically, for a batch of images denoted as B, our all-min contrastive loss is formulated as follows:

$$\mathcal{L}_{\text{all-min}} = -\frac{1}{|B|} \sum_{I_a \in B} \log \frac{\sum_{I_{b+} \in S^+(I_a)} \sum_{t=1}^{T} \exp(\sin(P_t^a, P_t^{b+})/\tau)}{\sum_{I_{b-} \in S^-(I_a)} \exp(\min_{t \in [1,T]} \sin(P_t^a, P_t^{b-})/\tau)},$$
(5)

where  $I_a$  represents an image in the batch B, the sim means the cosine similarity between corresponding parts from two different images or two different augmentation, the sets  $S^+(I_a)$  and  $S^-(I_a)$  represent the set of positive and negative image samples of  $I_a$ , respectively. If  $I_a$  is labeled, i.e.,  $I_a \in \mathcal{D}_l$ , the positive set  $S^+(I_a)$  consists of labeled images  $I_i \in \mathcal{D}_l$  whose class  $c_i$  is the same as the class  $c_a$  of image  $I_a$ , i.e.,  $S^+(I_a) = \{I_i \in \mathcal{D}_l : c_i = c_a\} \cap B$ . Otherwise, the set only contains augmentation images, i.e.,  $S^+(I_a) =$  $\{I_i \in B : ID(I_i) = ID(I_a)\}$ , where ID returns the id of the image. Contrary to the positive sets, if  $I_a$  is labeled, the negative set consists of labeled images with different classes from  $I_a$ , i.e.,  $S^-(I_a) = \{I_i \in \mathcal{D}_l : c_i \neq c_a\} \cap B$ . For unlabeled  $I_a$ , we randomly sample some images from B as negative samples, ensuring that not all parts of these images are similar to each corresponding part of  $I_a$ , and use the similarity of all positive pairs in the batch as the threshold for sampling negative images.

**Diverse Object Parts.** To prevent different part queries from collapsing onto the same semantic parts, we apply diversity loss to the part queries, encouraging them to focus on distinct object parts. Specifically, we constrain that different parts within an image are orthogonal:

$$\mathcal{L}_{\text{div}} = \sum_{P_t \in P_{part}} \sum_{\hat{P}_t \in P_{part} \setminus P_t} \max(0, \sin(P_t, \hat{P}_t)), \quad (6)$$

where  $P_t$  and  $\hat{P}_t$  are the part features of two different parts within an image, and sim denotes the cosine similarity.

#### 3.4. A Plug-and-Play Integration

Since our part discovery and learning do not rely on any specific GCD framework but only require a DINO backbone adopted by all existing GCD methods, we can integrate it into various GCD frameworks, such as the parameterized classification head in SimGCD [43] and SPTNet [42]. The parametrized framework freezes most layers of DINO and only tunes the final layer, treating the features of the [CLS] token,  $F_{token}$ , as the sample's feature representation for class prediction. Its learning objective includes both representation learning with InfoNCE loss and parametric classification on  $F_{token}$  [42, 43], formulated as follows:

$$\mathcal{L}_{\text{GCD}} = (1 - \lambda)(\mathcal{L}_{\text{cls}}^u + \mathcal{L}_{\text{rep}}^u) + \lambda(\mathcal{L}_{\text{cls}}^s + \mathcal{L}_{\text{rep}}^s) \quad (7)$$

where  $\lambda$  is the balance factor,  $\mathcal{L}_{cls}^{u} + \mathcal{L}_{rep}^{u}$  and  $\mathcal{L}_{cls}^{s} + \mathcal{L}_{rep}^{s}$  represent the sum of representation learning loss and parametric classification loss on labeled and unlabeled sets, respectively.

To solely showcase the discriminative yet generalizable effectiveness of our object parts, we simply pool the features from different parts to obtain an overall part feature, which replaces the feature  $F_{CLS}$  of CLS token for both classification prediction and feature representation learning, as illustrated in Figure 2b. The overall learning objective of the model is as follows:

$$\mathcal{L}_{\text{overall}} = (1 - \lambda)(\mathcal{L}_{\text{cls}}^u + \mathcal{L}_{\text{rep}}^u) + \lambda(\mathcal{L}_{\text{cls}}^s + \mathcal{L}_{\text{rep}}^s) + \mathcal{L}_{\text{div}} + \mathcal{L}_{\text{all-min}}$$
(8)

#### 4. Experiment

## 4.1. Experiment Setup

**Datasets.** We follow previous approaches [40, 43] to evaluate our method on four fine-grained and three generic image classification datasets. Following them, we adopt the same partitioning strategy across these datasets: for all fine-grained datasets [22, 27, 36, 41], CIFAR-10 [23], and ImageNet-100 [9], 50% of the classes are considered as known classes, with the remaining classes treated as novel classes. For CIFAR-100 [23], 80 classes are designated as known classes. Within these known classes, 50% samples from each class are chosen to construct the labeled set  $D_u$ , while the remaining samples constitute the unlabeled set  $D_u$ , along with the images of novel classes. For detailed information, please refer to the Appendix.

**Evaluation Metric.** We evaluate the effectiveness of our method on unlabeled dataset  $\mathcal{D}_u$  using clustering accuracy, following [7, 42, 43]. The accuracy is defined as ACC =  $\frac{1}{|\mathcal{D}_u|} \sum_{i=1}^{|\mathcal{D}_u|} 1(y_i = f^*(\hat{y}^i))$ , where  $f^* = \arg \max_{f \in \mathcal{F}} \sum_{i=1}^M 1(y_i = f(\hat{y}^i))$ , and  $\mathcal{F}$  is all possible matching between the predicted labels with the ground truth.

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

Table 1 compares our method with state-of-the-art models (SOTAs) on fine-grained datasets, including the SSB benchmark and the Herbarium19 dataset. Our method achieves state-of-the-art performance across multiple datasets and

Methods	CUB			Stanford Cars			FGVC-Aircraft			Herbarium-19		
moundus	All	Known	Novel	All	Known	Novel	All	Known	Novel	All	Known	Novel
k-means [2]	34.3	38.9	32.1	12.8	10.6	13.8	16.0	14.4	16.8	13.0	12.2	13.4
RS+ [15]	33.3	51.6	24.2	28.3	61.8	12.1	26.9	36.4	22.2	27.9	55.8	12.8
UNO+ [12]	35.1	49.0	28.1	35.5	70.5	18.6	40.3	56.4	32.2	28.3	53.7	14.7
ORCA [3]	35.3	45.6	30.2	23.5	50.1	10.7	22.0	31.8	17.1	20.9	30.9	15.5
GCD [40]	51.3	56.6	48.7	39.0	57.6	29.9	45.0	41.1	46.9	35.4	51.0	27.0
CiPR [16]	57.1	58.7	55.6	47.0	61.5	40.1	-	-	-	-	-	-
GPC [47]	52.0	55.5	47.5	38.2	58.9	27.4	43.3	40.7	44.8	-	-	-
DCCL [30]	63.5	60.8	64.9	43.1	55.7	36.2	-	-	-	-	-	-
PromptCAL [45]	62.9	64.4	62.1	50.2	70.1	40.6	52.2	52.2	52.3	37.0	52.0	28.9
PMI [6]	62.7	75.7	56.2	43.1	66.9	31.6	-	-	-	42.3	56.1	34.8
SimGCD [43]	60.3	65.6	57.7	53.8	71.9	45.0	54.2	59.1	51.8	43.0	58.0	35.1
Ours (SimGCD)	64.5	68.1	62.1	60.1	77.6	51.2	56.6	60.2	54.8	44.9	58.1	37.9
$\Delta$	+4.2	+2.5	+4.4	+6.3	+5.7	+6.2	+2.4	+1.1	+3.0	+1.9	+0.1	+2.8
SPTNet [42]	65.8	68.8	65.1	59.0	79.2	49.3	59.3	61.8	58.1	43.4	58.7	35.2
Ours (SPTNet)	68.5	73.1	66.2	62.3	80.7	53.4	60.9	63.5	59.6	47.3	66.6	37.0
$\Delta$	+2.7	+4.3	+1.1	+3.3	+1.5	+4.1	+1.6	+1.7	+1.5	+3.9	+7.9	+1.8
CMS [7]	68.2	76.5	64.0	56.9	76.1	47.6	56.0	63.4	52.3	36.4	54.9	26.4
$CMS^{\dagger}$	66.6	74.8	62.6	56.1	76.4	47.1	54.1	61.7	50.3	35.9	55.1	25.6
Ours (CMS <sup><math>\dagger</math></sup> )	69.2	75.3	66.1	62.1	81.8	52.6	57.9	66.0	53.8	38.4	57.2	28.2
$\Delta$	+2.6	+0.5	+3.5	+6.0	+5.4	+5.5	+3.8	+4.3	+3.5	+2.5	+2.1	+2.6

Table 1. Evaluation on the Semantic Shift Benchmark (SSB) and Herbarium-19 benchmarks.<sup>†</sup> represents our reproduced result.

consistently outperforms our three baselines, SimGCD [43] SPTNet [42] and CMS [7]. Specifically, compared to SO-TAs, we achieve average improvements of 2.2% on the SSB dataset across novel categories and notable enhancements on the more challenging Herbarium dataset, with improvements of 3.9%, 7.9%, and 1.8% across all, known, and novel categories, respectively.

Compared to the pioneering parameterized-head SimGCD, we achieve direct performance improvements on the four fine-grained datasets, with gains of 4.2%, 6.3%, 2.4%, and 1.9%, respectively. This indicates the effectiveness of our part discovery, as we simply replace the CLS token used by SimGCD in representation learning and class prediction with the average of our different part features. Particularly, we note that our improvements on novel classes are even more significant, with gains of 4.4%, 6.2%, 3.0%, 2.8%, respectively. This indicates that our method, by distinguishing between discriminative and transferable parts within similar categories and facilitating feature sharing among transferable parts, aids in the generalization from known to novel classes. Moreover, compared to SPTNet, which also utilizes local information, our method still obtains performance gains of 2.7%, 3.3%, 1.6%, and 3.9%, respectively, demonstrating the positive impact of our explicit object part discovery and utilization.

Methods	CIF	AR-10	CIFA	R-100	ImageNet-100		
	All	Novel	All	Novel	All	Novel	
InfoSieve [31]	94.8 05.0	93.4 04.8	78.3	70.5	80.5 83 1	73.8	
Ours(SimGCD)	93.9 97.1	94.8 98.2	82.3 80.9	70.3 <b>79.5</b>	83.2	77.8 78.5	

Table 2. Evaluation on generic image recognition datasets.

Unlike SPTNet, which implicitly enhances attention to local information through spatial visual prompts, our method directly extracts semantically meaningful object parts from images adaptively, which is more intuitive and interpretable. Additionally, by introducing our part learning loss on part tokens, we emphasize the diversity of parts and their adaptive discrimination between paired images, enhancing feature discriminability and generalization compared to globally unified visual prompts. In contrast to DCCL and InfoSieve, which top-down establish category relationships hierarchically or superclass-based, our approach establishes relationships from bottom to top through parts, resulting in enhanced generalization and boosting the accuracy of novel categories by 7.0% and 7.1% on the Stanford Cars and FGVC-Aircraft datasets, respectively.

Additionally, We compare our APL's effectiveness on a range of generic image recognition datasets with other methods which focus on fine-grained gneralized category discovery. Table 2 presents results on CIFAR-10, CIFAR-100 and ImageNet-100. Unlike fine-grained datasets, they require fewer fine-grained parts to distinguish between similar but different categories and share less knowledge across vastly different categories. Despite this, we consistently outperform all our baselines, with improvements of 3.4%, 3.2% and 0.7% on novel categories of CIFAR-10, CIFAR-100 and ImageNet-100, respectively. We present the t-SNE visualization of our method in Figure 4 and the comparison of similarity distribution of parts in Figure 5.

#### 4.3. Ablation Study

We employ SimGCD as the baseline and conduct ablation studies to validate the effectiveness of different components

in our APL. More ablations are shown in the Appendix.

Part Discovery. Table 3 showcases the results of different approaches and various variants of our method for part discovery. (1) Compared to the baseline SimGCD, the decreased performance observed (row 2) when clustering patches from individual images to discover parts indicates that naively incorporating local information can introduce noise. Moreover, single-image clustering struggles with consistency and lacks robustness. (2) Directly treating our extracted part priors  $F_{prior}$ , defined in Equ 3, as parts and using them to complement the CLS feature in SimGCD slightly improves its performance, demonstrating the positive impact of meaningful part information in GCD (row 3). (3) Moreover, we further introduce our learnable part queries to iteratively and dynamically query the DINO priors  $F_{\tt prior}$  and image patches  $F_{\tt patch}$  to address the consistency and correspondence challenges in GCD's part discovery (row 4). The significant performance improvements, i.e., 4.1% on novel categories, demonstrate the robustness of using learnable queries to capture cross-image shared part semantics. (4) We enhance the soft assignment of learnable queries with image patch features  $F_{patch}$  from the previous model to a hard assignment based on the Gumbel-Softmax, defined in Equ 4, which led to further improvements of 1.6%.

	Method	All	Known	Novel
1	SimGCD	53.8	71.9	45.0
2	1+ k-Means	51.0	70.3	41.7
3	1+ DINO part prior	55.1	73.9	46.0
4	3 +Learnable queries	58.5	74.7	50.1
5	4+ Hard assignment	60.1	77.6	51.2

Table 3. Ablation study of part discovery.

	Method	All	Known	Novel
1	SimGCD + part discovery	55.1	72.5	46.7
2	1+ part contrastive loss	55.2	73.2	46.5
3	1+ min-max contrastive loss	57.8	75.7	49.1
4	1+ diversity loss	56.7	72.0	49.6
5	3+ diversity loss	60.1	77.6	51.2

Table 4. Ablation study of learning objectives.

**Part Learning Objective.** Table 4 presents the performance of our method under various part learning objectives. (1) The baseline, developed on the SimGCD framework, integrates our part discovery, specifically  $P_{part}$ , as defined in Equ 4. However, it only incorporates the SimGCD loss  $\mathcal{L}_{GCD}$  from Equ 7 and does not employ any of our part learning losses. (2) Row 2 extends the baseline by implementing a common contrastive loss among parts, ensuring that all parts of positive samples are dissimilar. The results indicate



Figure 4. t-SNE visualization of part features on CUB-200 and image features on randomly sampled 10 classes on CIFAR-100. Each color represents a specific part (category).



Figure 5. The similarity distribution of discriminative parts between similar but distinct categories.

no improvement for novel categories but a slight enhancement in known categories, attributed to enhanced discriminability among fine-grained object parts across images. (3) When we substitute our all-min contrastive loss, we observe significant improvements of 2.7%, 2.8%, 2.4% in the all, known, and novel categories, respectively. It adaptively identifies the most discriminative part in negative samples to boost discriminability, while permitting similarity in other parts to enhance the transferability of shared parts across different categories. (4) We also evaluate the independent effect of diversity loss on our baseline, showing that emphasizing the distinctiveness of various parts within an image effectively enhances the discriminability of part-level features. (5) Finally, we validate the combined effects of the two part learning losses.

#### 5. Conclusion

This paper underscores the trade-off between discriminability and generalization in existing methods and introduces APL, which enhances both aspects simultaneously and seamlessly integrates into GCD frameworks. APL utilizes DINO prior and introduces novel all-min contrastive and diversity loss functions to enhance knowledge transfer, resulting in superior performance across benchmarks. **Acknowledgment:** This work was supported by the National Natural Science Foundation of China (No.62206174).

## References

- Ananthu Aniraj, Cassio F Dantas, Dino Ienco, and Diego Marcos. Pdiscoformer: Relaxing part discovery constraints with vision transformers. *arXiv preprint arXiv:2407.04538*, 2024. 3
- [2] David Arthur, Sergei Vassilvitskii, et al. k-means++: The advantages of careful seeding. In Soda, pages 1027–1035, 2007. 7
- Kaidi Cao, Maria Brbic, and Jure Leskovec. Open-world semi-supervised learning. arXiv preprint arXiv:2102.03526, 2021. 7
- [4] 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. 1, 3, 4
- [5] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: deep learning for interpretable image recognition. *Advances in neural information processing systems*, 32, 2019. 3
- [6] Florent Chiaroni, Jose Dolz, Ziko Imtiaz Masud, Amar Mitiche, and Ismail Ben Ayed. Parametric information maximization for generalized category discovery. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 1729–1739, 2023. 7
- [7] Sua Choi, Dahyun Kang, and Minsu Cho. Contrastive meanshift learning for generalized category discovery. arXiv preprint arXiv:2404.09451, 2024. 3, 6, 7
- [8] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Unsupervised part discovery from contrastive reconstruction. Advances in Neural Information Processing Systems, 34:28104–28118, 2021. 3, 4
- [9] 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. 1, 6
- [10] 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. 1
- [11] 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. 3, 4
- [12] 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
- [13] Emil Julius Gumbel. Statistical theory of extreme values and some practical applications: a series of lectures. US Government Printing Office, 1954. 5

- [14] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 3
- [15] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. arXiv preprint arXiv:2002.05714, 2020. 7
- [16] Shaozhe Hao, Kai Han, and Kwan-Yee K Wong. Cipr: An efficient framework with cross-instance positive relations for generalized category discovery. arXiv preprint arXiv:2304.06928, 2023. 7
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016. 4
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3
- [20] Zixuan Huang and Yin Li. Interpretable and accurate finegrained recognition via region grouping. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 8662–8672, 2020. 3
- [21] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 869–878, 2019. 3
- [22] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In Proceedings of the IEEE international conference on computer vision workshops, pages 554–561, 2013. 6
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [24] Xiangtai Li, Shilin Xu, Yibo Yang, Guangliang Cheng, Yunhai Tong, and Dacheng Tao. Panoptic-partformer: Learning a unified model for panoptic part segmentation. In *European Conference on Computer Vision*, pages 729–747. Springer, 2022. 3
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13, pages 740–755. Springer, 2014. 1
- [26] Chris J Maddison, Daniel Tarlow, and Tom Minka. A\* sampling. Advances in neural information processing systems, 27, 2014. 5
- [27] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classi-

fication of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 6

- [28] Umberto Michieli, Edoardo Borsato, Luca Rossi, and Pietro Zanuttigh. Gmnet: Graph matching network for large scale part semantic segmentation in the wild. In *Computer Vision– ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*, pages 397–414. Springer, 2020. 3
- [29] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193, 2023. 2, 3, 4
- [30] Nan Pu, Zhun Zhong, and Nicu Sebe. Dynamic conceptional contrastive learning for generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7579–7588, 2023. 1, 3, 7
- [31] Sarah Rastegar, Hazel Doughty, and Cees Snoek. Learn to categorize or categorize to learn? self-coding for generalized category discovery. Advances in Neural Information Processing Systems, 36, 2024. 1, 7
- [32] Sarah Rastegar, Mohammadreza Salehi, Yuki M Asano, Hazel Doughty, and Cees GM Snoek. Selex: Self-expertise in fine-grained generalized category discovery. arXiv preprint arXiv:2408.14371, 2024. 7
- [33] Oindrila Saha and Subhransu Maji. Particle: Part discovery and contrastive learning for fine-grained recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 167–176, 2023. 3
- [34] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 1
- [35] Peize Sun, Shoufa Chen, Chenchen Zhu, Fanyi Xiao, Ping Luo, Saining Xie, and Zhicheng Yan. Going denser with open-vocabulary part segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15453–15465, 2023. 3
- [36] Kiat Chuan Tan, Yulong Liu, Barbara Ambrose, Melissa Tulig, and Serge Belongie. The herbarium challenge 2019 dataset. arXiv preprint arXiv:1906.05372, 2019. 6
- [37] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16, pages 776–794. Springer, 2020. 3
- [38] Robert van der Klis, Stephan Alaniz, Massimiliano Mancini, Cassio F Dantas, Dino Ienco, Zeynep Akata, and Diego Marcos. Pdisconet: Semantically consistent part discovery for fine-grained recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1866– 1876, 2023. 3
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. Advances in neural information processing systems, 30, 2017. 3
- [40] 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. 1, 3, 6, 7

- [41] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 6
- [42] Hongjun Wang, Sagar Vaze, and Kai Han. Sptnet: An efficient alternative framework for generalized category discovery with spatial prompt tuning. *arXiv preprint arXiv:2403.13684*, 2024. 1, 3, 4, 6, 7
- [43] Xin Wen, Bingchen Zhao, and Xiaojuan Qi. Parametric classification for generalized category discovery: A baseline study. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16590–16600, 2023. 1, 2, 3, 4, 6, 7
- [44] Tingyu Weng, Jun Xiao, and Haiyong Jiang. Decompose novel into known: Part concept learning for 3d novel class discovery. Advances in Neural Information Processing Systems, 36:54282–54299, 2023. 3
- [45] Sheng Zhang, Salman Khan, Zhiqiang Shen, Muzammal Naseer, Guangyi Chen, and Fahad Shahbaz Khan. 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, pages 3479–3488, 2023. 1, 7
- [46] 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. 3
- [47] Bingchen Zhao, Xin Wen, and Kai Han. Learning semisupervised gaussian mixture models for generalized category discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16623–16633, 2023.
   7
- [48] Yifan Zhao, Jia Li, Yu Zhang, and Yonghong Tian. Multiclass part parsing with joint boundary-semantic awareness. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 9177–9186, 2019. 3
- [49] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for finegrained image recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 5209–5217, 2017. 3