

Generalized Class Discovery in Instance Segmentation

Cuong Manh Hoang¹, Yeejin Lee¹, Byeongkeun Kang^{2*}

¹Seoul National University of Science and Technology, Republic of Korea

²Chung-Ang University, Republic of Korea

{cuonghoang, yeejinlee}@seoultech.ac.kr, byeongkeunkang@cau.ac.kr

Abstract

This work addresses the task of generalized class discovery (GCD) in instance segmentation. The goal is to discover novel classes and obtain a model capable of segmenting instances of both known and novel categories, given labeled and unlabeled data. Since the real world contains numerous objects with long-tailed distributions, the instance distribution for each class is inherently imbalanced. To address the imbalanced distributions, we propose an instance-wise temperature assignment (ITA) method for contrastive learning and class-wise reliability criteria for pseudo-labels. The ITA method relaxes instance discrimination for samples belonging to head classes to enhance GCD. The reliability criteria are to avoid excluding most pseudo-labels for tail classes when training an instance segmentation network using pseudo-labels from GCD. Additionally, we propose dynamically adjusting the criteria to leverage diverse samples in the early stages while relying only on reliable pseudo-labels in the later stages. We also introduce an efficient soft attention module to encode object-specific representations for GCD. Finally, we evaluate our proposed method by conducting experiments on two settings: COCO_{half} + LVIS and LVIS + Visual Genome. The experimental results demonstrate that the proposed method outperforms previous state-of-the-art methods.

1 Introduction

While supervised instance segmentation methods (He et al. 2017; Cheng et al. 2022; Jain et al. 2023) have achieved impressive performance, they require large-scale datasets with expensive human annotations. To reduce annotation costs, researchers have investigated semi-supervised learning (SSL) methods (Bellver et al. 2019; Yang et al. 2023; Berrada et al. 2024) that utilize unlabeled images along with small-scale labeled data, as well as weakly-supervised methods (Lan et al. 2021) relying on weak annotations. However, all these methods rely on the closed-world assumption and can recognize only the objects belonging to the classes (*i.e.*, known classes) in the labeled dataset.

To address this limitation, researchers have introduced novel category discovery (NCD) (Han, Vedaldi, and Zisserman 2019). Unlike SSL, where the unlabeled images contain only known classes, NCD assumes that the unlabeled

data include novel categories. Recently, generalized (novel) category discovery (GCD) (Vaze et al. 2022; Cao, Brbic, and Leskovec 2022) was introduced, further relaxing the assumptions on the unlabeled data. It assumes that the unlabeled data may contain both known and novel classes, making the problem more challenging and realistic. Given labeled and unlabeled data, GCD aims to train a model capable of recognizing both the known classes (*e.g.*, person and car) in the labeled data and the novel categories (*e.g.*, unknown₁ and unknown₂) discovered from the unlabeled data.

Most previous works have investigated GCD for curated and balanced image classification datasets (Vaze et al. 2022; An et al. 2023; Zhang et al. 2023; Pu, Zhong, and Sebe 2023; Wen, Zhao, and Qi 2023). Recently, researchers have explored GCD for image classification on imbalanced datasets (Bai et al. 2023; Li et al. 2023; Li, Meinel, and Yang 2023), semantic segmentation (Zhao et al. 2022), 3D point cloud semantic segmentation (Riz et al. 2023), and instance segmentation (Fomenko et al. 2022; Weng et al. 2021). Most of these works leverage semi-supervised contrastive learning and pseudo-label generation regardless of tasks.

In this work, we also investigate GCD in instance segmentation. It is worth noting that because the real world contains numerous objects with long-tailed distributions, the instance distribution for each class in instance segmentation datasets is inherently imbalanced. For example, a ‘car’ appears more frequently than an ‘ashtray’. (1) To address this imbalanced distribution, we propose an instance-wise temperature assignment method for contrastive learning. While typical contrastive learning losses treat samples from head and tail classes equally, we aim to emphasize group-wise discrimination for head class samples while focusing on instance-wise discrimination for tail class samples, inspired by (Kukleva et al. 2023). (2) Although relying on reliable pseudo-labels is important (Yang et al. 2022), applying fixed and global reliability criteria to imbalanced data tends to exclude most pseudo-labels for tail classes. Therefore, we propose to utilize class-wise reliability criteria to apply varying thresholds for head and tail classes. Additionally, we dynamically adjust the reliability criteria throughout training to leverage diverse samples in the early stages while focusing on reliable pseudo-labels in later stages. (3) Finally, we introduce an efficient soft attention module based on spatial pooling and depth reduction to effectively encode repre-

*Corresponding author.

sentations for target instances while suppressing those from background or adjacent objects.

The contributions of this paper are as follows: (1) We propose an instance-wise temperature assignment method for semi-supervised contrastive learning in GCD to enhance the separability of classes in long-tailed distributions. (2) We introduce a reliability-based dynamic learning method for training with pseudo-labels to apply different reliability criteria to each class based on its tailness. Additionally, we adjust these criteria during training to rely on strictly reliable pseudo-labels in the later stages while using diverse data in the early stages. (3) We propose an efficient soft attention module for encoding object-specific representations for GCD. (4) We validate the effectiveness of the proposed method by discovering novel classes and training an instance segmentation network using labeled and unlabeled data.

2 Related Works

Generalized Class Discovery in Image Classification.

(Vaze et al. 2022) introduced GCD, which aims to categorize unlabeled images given both labeled and unlabeled data. They trained an embedding network using unsupervised contrastive learning on all the data and supervised contrastive learning on the labeled data. They then applied semi-supervised k -means clustering to assign class or cluster labels to the unlabeled images. Additionally, they proposed a method for estimating the number of novel classes.

(An et al. 2023) proposed DPN, which utilizes two sets of category-wise prototypes: one for labeled data and the other for unlabeled images based on k -means clustering. For clustering, they assumed prior knowledge of the total number of categories. They discovered novel classes in the unlabeled data by applying the Hungarian algorithm (Kuhn 1955) to the two sets. (Zhang et al. 2023) presented Prompt-CAL, which uses an affinity graph to generate pseudo-labels for unlabeled data. (Pu, Zhong, and Sebe 2023) introduced the DCCL framework, which employs a hyperparameter-free clustering algorithm (Rosvall and Bergstrom 2008) to generate pseudo-labels in the absence of ground-truth cluster numbers. (Wen, Zhao, and Qi 2023) proposed SimGCD, a one-stage framework that replaces the separate semi-supervised clustering in (Vaze et al. 2022) with a jointly trainable parametric classifier. They analyzed the problem of using unreliable pseudo-labels for training a parametric classifier and proposed using soft pseudo-labels.

Recently, (Bai et al. 2023) introduced GCD for long-tailed datasets to address imbalanced distributions in real-world scenarios. They proposed the BaCon framework, which includes a pseudo-labeling branch and a contrastive learning branch. To handle imbalanced distributions, they estimated data distribution using k -means clustering and the Hungarian algorithm (Kuhn 1955). Concurrently, (Li et al. 2023) proposed ImbaGCD, which estimates class prior distributions assuming unknown classes are usually tail classes. They then generated pseudo-labels for unlabeled images using the estimated prior distribution and the Sinkhorn-Knopp algorithm (Cuturi 2013). Later, (Li, Meinel, and Yang 2023) replaced the expectation-maximization (EM) and Sinkhorn-Knopp (Cuturi 2013) algorithms in (Li et al. 2023) with

cross-entropy-based regularization losses to reduce computational costs.

Class Discovery in Segmentation and Detection. (Zhao et al. 2022) introduced NCD in semantic segmentation. They proposed to find novel salient object regions using a saliency model and a segmentation network trained on labeled data. They then applied clustering to these object regions to obtain pseudo-labels. Finally, they trained a segmentation network using the labeled data and the unlabeled images with clean pseudo-labels from clustering and online pseudo-labels. The clean pseudo-labels were dynamically assigned based on entropy ranking. (Riz et al. 2023) extended the method from (Zhao et al. 2022) to 3D point cloud semantic segmentation. Specifically, they utilized online clustering and exploited uncertainty quantification to generate pseudo-labels.

In instance segmentation, (Weng et al. 2021) investigated unsupervised long-tail category discovery. They initially employed a class-agnostic mask proposal network to obtain masks for all objects. They then trained an embedding network using self-supervised triplet losses. Finally, they applied hyperbolic k -means clustering to discover novel categories. (Fomenko et al. 2022) introduced novel class discovery and localization, which can be viewed as GCD in object detection and instance segmentation. They first trained a Faster R-CNN (Ren et al. 2015) or Mask R-CNN (He et al. 2017) using the labeled data and froze the network except for its classification head. Then, they applied the frozen network to unlabeled and labeled images to obtain region proposals. Subsequently, they expanded the classification head to incorporate new classes and trained it using pseudo-labels generated by online clustering based on the Sinkhorn-Knopp algorithm (Cuturi 2013).

3 Proposed Method

We first define the GCD problem in instance segmentation, which aims to discover novel classes and learn to segment instances of both known and novel categories, given labeled and unlabeled data. We then present our GCD method in Section 3.2 and the method for training an instance segmentation network in Section 3.3. An overview of the proposed framework during training is illustrated in Figure 1.

3.1 Preliminaries

Problem Formulation. We are given a labeled dataset \mathcal{D}^l and an unlabeled dataset \mathcal{D}^u . \mathcal{D}^l contains images $\{\mathbf{I}^l\}$ along with instance-wise class and mask labels ($\{\mathbf{y}^l\}$, $\{\mathbf{M}^l\}$) for known classes \mathcal{C}^k , while \mathcal{D}^u comprises only images $\{\mathbf{I}^u\}$. Given \mathcal{D}^l and \mathcal{D}^u , GCD in instance segmentation aims to discover novel categories \mathcal{C}^n (i.e., $\mathcal{C}^k \cap \mathcal{C}^n = \emptyset$) and to obtain a model capable of segmenting instances of both the known and novel classes $\mathcal{C} = \mathcal{C}^k \cup \mathcal{C}^n$. Hence, during inference, the network is expected to segment instances of known classes (e.g., person and car) as well as novel categories (e.g., unknown₁ and unknown₂) given an image \mathbf{I} . The images in \mathcal{D}^l and \mathcal{D}^u may contain instances of both the known and novel classes.

Contrastive Learning for GCD. (Vaze et al. 2022) introduced a contrastive learning (CL) method for GCD in bal-

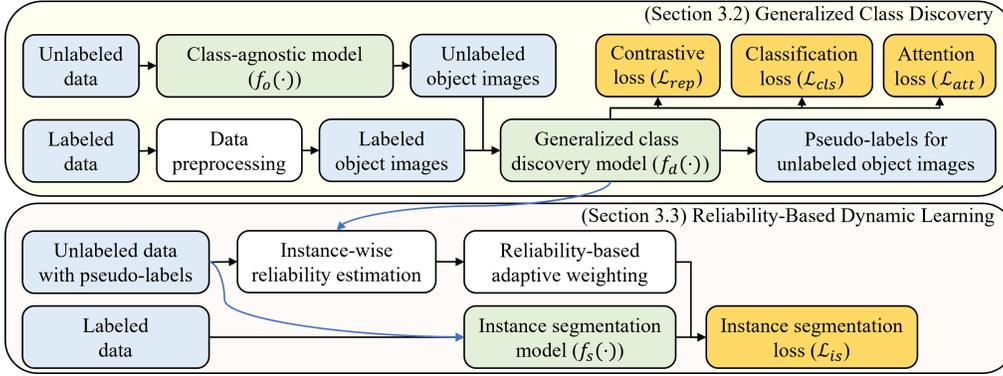


Figure 1: Overview of the proposed framework during training. We first train a class-agnostic instance segmentation network $f_o(\cdot)$ and apply it to unlabeled images to generate class-agnostic instance masks. Then, we train the GCD model $f_d(\cdot)$ using unlabeled object images and labeled object images to discover novel classes in the unlabeled data. Finally, we train an instance segmentation network $f_s(\cdot)$ using the labeled data and the unlabeled images with pseudo-labels.

anced image classification datasets. They first pre-trained a backbone with DINO (Caron et al. 2021) on the ImageNet dataset (Russakovsky et al. 2015) without labels. Subsequently, they fine-tuned the backbone and a projection head using supervised CL on the labeled data and unsupervised CL on both the labeled and unlabeled data.

Following (Wen, Zhao, and Qi 2023; Vaze et al. 2022), our GCD model $f_d(\cdot)$ consists of a backbone $b(\cdot)$ and a projection head $g(\cdot)$. We utilize an MLP for $g(\cdot)$ and a ResNet-50 backbone (He et al. 2016) for $b(\cdot)$, which is pre-trained on the unlabeled ImageNet dataset (Russakovsky et al. 2015) using DINO (Caron et al. 2021). Additionally, we employ a momentum encoder $f'_d(\cdot)$ whose parameters are momentum-based moving averages of the parameters of $f_d(\cdot)$ during training, following MoCo (He et al. 2020). We also use a queue to store the embeddings from $f'_d(\cdot)$ for the samples of both the previous and current mini-batches.

Formally, given an image, we generate two views (random augmentations) I_i and I'_i . We then encode them using $f_d(\cdot)$ and $f'_d(\cdot)$ to obtain $z_i = f_d(I_i) = g(b(I_i))$ and $z'_i = f'_d(I'_i)$, respectively. We store z'_i from the samples in the previous and current mini-batches in the queue \mathcal{Z}' . In (Vaze et al. 2022), the unsupervised contrastive loss \mathcal{L}_{rep}^u and supervised contrastive loss \mathcal{L}_{rep}^s are computed as follows:

$$\begin{aligned} \mathcal{L}_{rep}^u &:= -\log \frac{\exp(z_i^\top z'_i / \tau)}{\sum_{z'_j \in \hat{\mathcal{Z}}_i} \exp(z_i^\top z'_j / \tau)}, \\ \mathcal{L}_{rep}^s &:= -\frac{1}{|\mathcal{Z}_i^p|} \sum_{z'_k \in \mathcal{Z}_i^p} \log \frac{\exp(z_i^\top z'_k / \tau)}{\sum_{z'_j \in \hat{\mathcal{Z}}_i} \exp(z_i^\top z'_j / \tau)} \end{aligned} \quad (1)$$

where $\hat{\mathcal{Z}}_i$ represents the set \mathcal{Z}' excluding z'_i (i.e., $\hat{\mathcal{Z}}_i = \mathcal{Z}' \setminus z'_i$); \mathcal{Z}_i^p denotes the subset of \mathcal{Z}' containing the representations that belong to the same class as I_i ; τ is a temperature hyperparameter; and $|\cdot|$ denotes the number of samples in the set.

3.2 Generalized Class Discovery

Given \mathcal{D}^l and \mathcal{D}^u , we aim to discover novel categories in \mathcal{D}^u using the knowledge from \mathcal{D}^l . To achieve this, we first train an instance segmentation network using \mathcal{D}^l and \mathcal{D}^u to generate class-agnostic instance masks M^u for all objects in \mathcal{D}^u . Subsequently, we crop the unlabeled images using M^u and the labeled images using the ground-truth masks M^l . Then, we train a GCD model using the cropped unlabeled images \mathcal{I}_o^u and the cropped labeled images \mathcal{I}_o^l with class labels \mathbf{y}^l to generate pseudo-class labels \mathbf{y}^u for \mathcal{I}_o^u .

Class-Agnostic Instance Mask Generation. Similar to (Fomenko et al. 2022), we first train an instance segmentation network $f_o(\cdot)$ to obtain class-agnostic instance masks for both known and unseen classes $\mathcal{C} = \mathcal{C}^k \cup \mathcal{C}^n$. We train a class-agnostic instance segmentation network, the Generic Grouping Network (GGN) from (Wang et al. 2022), using both \mathcal{D}^l and \mathcal{D}^u . We experimentally demonstrate that our GCD method is robust when applied to other class-agnostic instance segmentation methods.

Once the training terminates, we apply $f_o(\cdot)$ to the images $I^u \in \mathcal{D}^u$ to obtain instance masks M^u . We then construct an unlabeled object image set \mathcal{I}_o^u by cropping the rectangular regions of I^u based on M^u . Similarly, a labeled object image set \mathcal{I}_o^l is prepared by cropping I^l in \mathcal{D}^l based on the mask labels M^l .

Contrastive Learning for GCD in Instance Segmentation. We propose a contrastive learning method for GCD in instance segmentation by modifying the losses in Eq. (1). While these losses are designed for curated and balanced data, instances in typical instance segmentation datasets are naturally imbalanced (i.e., certain objects appear more frequently than others). To address the long-tail distribution of instances for each class, we propose to adjust the temperature parameters in \mathcal{L}_{rep}^u and \mathcal{L}_{rep}^s for each instance based on its likelihood of belonging to head classes.

Figure 2 visualizes t -SNE projections of two semantically similar classes: ‘book’ and ‘booklet’. Previously, (Kukleva

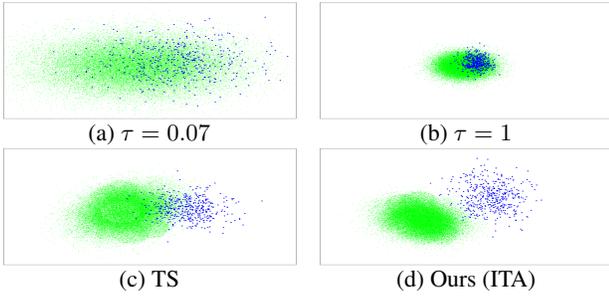


Figure 2: t -SNE visualization of two semantically close classes. Green: single head class (‘book’); Blue: single tail class (‘booklet’).

et al. 2023) showed that a low temperature value tends to uniformly discriminate all instances while a high temperature value leads to group-wise discrimination, as shown in Figure 2 (a) and (b). Based on this, (Kukleva et al. 2023) introduced a cosine temperature scheduling (TS) method to alternate between instance-wise and group-wise discrimination. In contrast, we propose to estimate the headness of each instance and assign high/low temperature values to instances belonging to head/tail class samples. Figure 2 demonstrates the superiority of our instance-wise temperature assignment (ITA) method compared to static assignments and TS (Kukleva et al. 2023).

In detail, we first compute a headness score \hat{h}_i for each instance I_i by estimating the density of the neighborhood of z_i in the embedding space. A higher score \hat{h}_i indicates a higher probability of I_i belonging to a head class. We then apply a momentum update to the headness scores to enhance the robustness of the estimation. Specifically, \hat{h}_i and the momentum-updated headness score h_i at the t -th epoch are computed as follows:

$$\hat{h}_i^t := \frac{\sum_{z'_j \in \hat{\mathcal{Z}}_i^{topK} \exp(z_i^T z'_j)}{\sum_{z'_j \in \hat{\mathcal{Z}}_i \exp(z_i^T z'_j)}, \quad (2)$$

$$h_i^t := \rho h_i^{t-1} + (1 - \rho) \hat{h}_i^t$$

where $\hat{\mathcal{Z}}_i = \mathcal{Z}' \setminus z'_i$; $\hat{\mathcal{Z}}_i^{topK}$ denotes the set containing the $K\%$ most similar representations to z_i in $\hat{\mathcal{Z}}_i$; ρ denotes a momentum hyperparameter with a value between 0 and 1.

Subsequently, we determine a temperature value τ_i for each instance I_i using h_i^t . To avoid extreme values, we constrain h_i^t to fall within the lowest 10% (h^{low}) and highest 10% (h^{high}) of the scores. We then apply min-max normalization to adjust the score to fall within the range between τ^{min} and τ^{max} . Specifically, the temperature value τ_i for I_i is calculated as follows:

$$\tau_i := \frac{\bar{h}_i^t - \min(\mathcal{H}^t)}{\max(\mathcal{H}^t) - \min(\mathcal{H}^t)} (\tau^{max} - \tau^{min}) + \tau^{min} \quad (3)$$

where $\bar{h}_i^t := \min(\max(h_i^t, h^{low}), h^{high})$; \mathcal{H}^t represents the set containing the headness scores for all samples. For efficiency, τ_i and h_i^t are updated at every epoch.

The two contrastive losses \mathcal{L}_{rep}^u and \mathcal{L}_{rep}^s in Eq. (1) are modified using the estimated instance-wise temperature value τ_i as follows:

$$\mathcal{L}_{rep}^u := -\log \frac{\exp(z_i^T z'_i / \tau_i)}{\sum_{z'_j \in \hat{\mathcal{Z}}_i \exp(z_i^T z'_j / \tau_i)}, \quad (4)$$

$$\mathcal{L}_{rep}^s := -\frac{1}{|\mathcal{Z}_i^p|} \sum_{z'_k \in \mathcal{Z}_i^p} \log \frac{\exp(z_i^T z'_k / \tau_i)}{\sum_{z'_j \in \hat{\mathcal{Z}}_i \exp(z_i^T z'_j / \tau_i)}$$

Soft Attention Module. Since \mathcal{I}_o^u and \mathcal{I}_o^l contain target objects along with background or adjacent objects, we investigate a soft attention module (SAM) to encode object-specific features. Although we generate pseudo-masks M^u for \mathcal{D}^u using $f_o(\cdot)$, directly using these pseudo-masks to encode object-specific features is risky due to noisy boundaries. To address this, we train an efficient attention module using the pseudo-masks M^u for \mathcal{D}^u and ground-truth masks M^l for \mathcal{D}^l . We integrate this attention module into every stage of the CNN backbone. Additionally, we utilize pooled feature maps and embedding functions with depth reduction to reduce computational complexity.

Given a feature map $F \in \mathbb{R}^{D \times \bar{H} \times \bar{W}}$ at the end of each stage in the backbone, we first reduce its dimensions to decrease subsequent computations. Specifically, we apply spatial average pooling to F with varying receptive fields. Then, we use M embedding functions $\eta_i(\cdot)$ to generate M outputs $P_i \in \mathbb{R}^{d \times s_i \times s_i}$. Here, i indexes the M outputs, d represents the reduced depth produced by the embedding functions, and $s_i \times s_i$ denotes the resulting spatial dimension after pooling. Subsequently, we reshape P_i into $\hat{P}_i \in \mathbb{R}^{d \times s_i^2}$ and concatenate them to obtain $\bar{P} \in \mathbb{R}^{d \times (s_1^2 + s_2^2 + \dots + s_M^2)}$.

Then, we compute a pairwise affinity matrix A by projecting F using an embedding function $\phi(\cdot)$ and multiplying the result by \bar{P} (i.e., $A := \bar{P}^T \phi(F)$). The matrix $A \in \mathbb{R}^{(s_1^2 + s_2^2 + \dots + s_M^2) \times \bar{H} \times \bar{W}}$ represents the spatial relations between the pooled feature map P and F . Additionally, we project F using a function $\psi(\cdot)$ and apply global average pooling along the channel dimension, generating a map G .

Finally, we concatenate A with G and apply an embedding function $\nu(\cdot)$ followed by a sigmoid function to obtain an attention map $S \in \mathbb{R}^{1 \times \bar{H} \times \bar{W}}$. This map S is then element-wise multiplied with each channel of F to produce the output $O \in \mathbb{R}^{D \times \bar{H} \times \bar{W}}$ of the attention module.

$$O := S \odot F := \sigma(\nu([A, G])) \odot F \quad (5)$$

where \odot and $\sigma(\cdot)$ denote element-wise multiplication and the sigmoid function, respectively; $[\cdot, \cdot]$ represents concatenation. Each of the embedding functions ($\eta(\cdot)$, $\psi(\cdot)$, $\phi(\cdot)$) consists of a 1×1 convolution layer, batch normalization, and ReLU activation. In comparison, $\nu(\cdot)$ contains only a 1×1 convolution layer and batch normalization.

To train the soft attention modules, we use object masks M^u from $f_o(\cdot)$ for \mathcal{D}^u and ground-truth masks M^l for \mathcal{D}^l . Because the pseudo-masks for \mathcal{D}^u are noisy, especially near object boundaries (Wang, Li, and Wang 2022), we utilize a

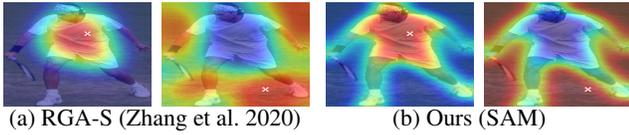


Figure 3: Visualization of the pairwise affinity between the white cross marked location and other pixels.

weight map \mathbf{W} to reduce reliance on these regions. Specifically, the attention loss \mathcal{L}_{att} is computed as follows:

$$\mathcal{L}_{att} := \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \mathbf{W}_{ij} \|\mathbf{S}_{ij} - \mathbf{M}_{ij}\|_2^2, \quad (6)$$

where \mathbf{W}_{ij} is set to w if $d_{ij} \leq \bar{d}$ and $\mathbf{M} \in \{\mathcal{M}^u\}$, and to 1 otherwise. Here, d_{ij} denotes the Euclidean distance from (i, j) to the nearest object boundary; \bar{d} is a hyperparameter that defines the boundary regions; and w is a weighting coefficient (≤ 1). Since SAM is applied after each stage of the backbone, \mathcal{L}_{att} is obtained by averaging all the corresponding losses.

Figure 3 visualizes the pairwise affinity between a marked position and other pixels. The results show that our method, which uses pooled feature maps, is more robust than the previous approach (Zhang et al. 2020).

Deep Clustering for GCD. To avoid the separate semi-supervised clustering step in (Vaze et al. 2022), we employ a deep clustering method similar to those in (Fomenko et al. 2022; Wen, Zhao, and Qi 2023). However, unlike (Fomenko et al. 2022), our method does not rely on an experimentally selected target prior distribution. Additionally, it is designed to handle imbalanced data, in contrast to (Wen, Zhao, and Qi 2023). Specifically, we use the method from (Zhang et al. 2021) for clustering \mathcal{I}_o^u , with a minor modification: replacing L2 distance with cosine similarity.

We compute the KL-divergence-based loss \mathcal{L}_{cls}^u on \mathcal{I}_o^u for unsupervised clustering in (Zhang et al. 2021), as follows:

$$\mathcal{L}_{cls}^u := \sum_{c=1}^C \bar{p}_{ic} \log \frac{\bar{p}_{ic}}{q_{ic}} \quad (7)$$

where q_{ic} is the probability of \mathbf{I}_i belonging to class c ; \bar{p}_{ic} is the auxiliary target class probability; and C is the total number of clusters/classes.

Independent from deep clustering, we additionally compute the typical cross-entropy loss \mathcal{L}_{cls}^s on \mathcal{D}^l for supervised classification, as follows:

$$\mathcal{L}_{cls}^s := \sum_{c=1}^C -y_{ic} \log q_{ic} \quad (8)$$

where \mathbf{y}_i denotes the one-hot encoded label for \mathbf{I}_i .

Total Loss for GCD. The total loss \mathcal{L}_{gcd} for $f_d(\cdot)$ is computed as the weighted sum of the two contrastive losses, the two classification losses, and the attention loss, as follows:

$$\mathcal{L}_{gcd} := \mathcal{L}_{att} + (1 - \lambda)\mathcal{L}_{rep}^u + \lambda\mathcal{L}_{rep}^s + (1 - \lambda)\mathcal{L}_{cls}^u + \lambda\mathcal{L}_{cls}^s \quad (9)$$

where λ is a hyperparameter used to balance the losses, following (Wen, Zhao, and Qi 2023).

3.3 Reliability-Based Dynamic Learning

We generate pseudo-masks \mathcal{M}^u and pseudo-class labels \mathbf{y}^u for \mathcal{D}^u using the method described in Section 3.2. Subsequently, we train an instance segmentation network $f_s(\cdot)$ that can segment instances of both known and novel classes using \mathcal{D}^l and \mathcal{D}^u with pseudo-labels. To address the issues of inaccurate pseudo-labels and imbalanced instance distributions across classes, we propose a reliability-based dynamic learning (RDL) method. It applies different reliability criteria to each class to avoid excluding all samples from tail classes. Additionally, it adjusts these criteria during training to use diverse data in the early stages while relying only on reliable pseudo-labels in the later stages.

Inspired by (Yang et al. 2022), we use holistic stability to measure the reliability of the pseudo-labels. At every fixed number of epochs during training $f_d(\cdot)$, we save the model at that point in time. We then apply these saved models to object images $\mathbf{I}_i \in \mathcal{I}_o^u$ to compute the probability q_{ic}^t of \mathbf{I}_i belonging to class c , where t denotes the index of the stored models, ranging from 1 to \bar{T} . Subsequently, we compute the stability s_i of the probabilities by comparing $q_{ic}^{\bar{T}}$ from the final model with q_{ic}^t from the intermediate models, as follows:

$$s_i = \sum_{t=1}^{\bar{T}-1} \frac{1}{KL(q_{ic}^{\bar{T}} \| q_{ic}^t)} \quad (10)$$

where $KL(\cdot \| \cdot)$ represents the Kullback-Leibler divergence, and $q_{ic}^t \in \mathbb{R}^{|\mathcal{C}|}$. Since a higher s_i indicates greater stability, (Yang et al. 2022) considered pseudo-labels with the lowest $r\%$ scores unreliable.

However, applying the same criteria to all data may result in categorizing most samples of a certain class as unreliable. Specifically, for imbalanced data, instances of tail classes tend to have lower s_i than those of head classes due to the smaller number of training samples. Additionally, because neural networks tend to first memorize easy samples and then gradually learn harder instances during training (Arpit et al. 2017), difficult samples/classes often have lower s_i than easier ones. To address these issues, we propose to use class-wise reliability criteria, which consider pseudo-labels with the lowest $r\%$ scores per class as unreliable.

Additionally, we gradually increase the portion $r\%$ of unreliable samples to initially learn from all data and later optimize using only reliable samples. The idea is that, in the early stages, having a larger number of diverse samples is more important than the accuracy of pseudo-labels, while pseudo-label quality becomes more crucial in later stages. Specifically, we first compute \bar{t}_i for each instance i by finding the rank of s_i among the lowest values within its class. We assign γ to \bar{t}_i if its proportional rank falls between $\frac{\gamma}{T_{is}}$ and $\frac{\gamma+1}{T_{is}}$. Then, the reliability-based adjustment weight κ_i^t is computed as follows:

$$\kappa_i^t = \sqrt{1 - \left(\max \left(\frac{t - \bar{t}_i}{T_{is}}, 0 \right) \right)^2} \quad (11)$$

where t and T_{is} denote the current epoch and the total number of epochs for training $f_s(\cdot)$, respectively.

Finally, we employ SOLOv2 (Wang et al. 2020) for $f_s(\cdot)$ and use its loss function with modifications for training. First, we replace the focal loss (Lin et al. 2017) with the equalized focal loss (Li et al. 2022), which performs better on imbalanced data. We use this modified loss \mathcal{L}_s for \mathcal{D}^l and this loss multiplied by κ_i^t for \mathcal{D}^u . Therefore, the total instance segmentation loss \mathcal{L}_{is} is computed as follows:

$$\mathcal{L}_{is} = \sum_{i \in \mathcal{B}^l} \mathcal{L}_s(I_i^l, M_i^l, \mathbf{y}_i^l) + \sum_{i \in \mathcal{B}^u} \kappa_i^t \mathcal{L}_s(I_i^u, M_i^u, \mathbf{y}_i^u) \quad (12)$$

where \mathcal{B}^l and \mathcal{B}^u denote the sets containing the indices of the data from \mathcal{D}^l and \mathcal{D}^u , respectively.

4 Experiments and Results

4.1 Experimental Setting

Dataset. We conducted experiments in two settings: COCO_{half} + LVIS and LVIS + VG, following (Fomenko et al. 2022). In the COCO_{half} + LVIS setting, we consider 80 COCO classes as known classes and aim to discover 1,123 disjoint LVIS classes from the total 1,203 LVIS classes. Among the 100K training images in the LVIS dataset (Gupta, Dollár, and Girshick 2019), we use 50K images with labels for the 80 COCO classes for \mathcal{D}^l and the entire 100K images without labels for \mathcal{D}^u . We use the 20K LVIS validation images for evaluation.

In the LVIS + VG setting, we utilize 1,203 LVIS classes as known classes and aim to discover 2,726 disjoint classes in the Visual Genome (VG) v1.4 dataset (Krishna et al. 2017). We use the entire 100K LVIS training data for \mathcal{D}^l and the combined 158K LVIS and VG training images for \mathcal{D}^u . For evaluation, we use 8K images that appear in both the LVIS and VG validation sets. Although the VG dataset contains over 7K classes, only 3,367 classes appear in both \mathcal{D}^u and the validation data. After excluding the 641 classes that overlap with the known classes, we aim to discover 2,726 classes.

Implementation Details. We use ResNet-50 as the backbone for both $f_d(\cdot)$ and $f_s(\cdot)$. For $f_d(\cdot)$, we apply the SAM to all four stages of the backbone. We train $f_d(\cdot)$ for 390 epochs with $K = 1\%$, $\tau^{min} = 0.07$, $\tau^{max} = 1$, and $\lambda = 0.35$. For the SAM, we set $M = 3$, $d = D/8$, $w = 0.25 \cdot stg$, $\hat{d} = 1$, $s_1 = \lfloor 18/(2^{stg-1}) \rfloor$, $s_2 = \lfloor 12/(2^{stg-1}) \rfloor$, and $s_3 = \lfloor 8/(2^{stg-1}) \rfloor$, where stg is the stage index and $\lfloor \cdot \rfloor$ denotes rounding. We train $f_s(\cdot)$ for 36 epochs with $\bar{T} = 3$. The experiments were conducted on a computer with two Nvidia GeForce RTX 3090 GPUs, an Intel Core i9-10940X CPU, and 128 GB RAM. The code will be publicly available on GitHub upon publication to ensure reproducibility.

Evaluation Metric. We first apply the Hungarian algorithm (Kuhn 1955) to find a one-to-one mapping between the discovered classes and the ground-truth novel classes, following (Fomenko et al. 2022). We then calculate $mAP_{.50:.05:.95}$ for all, known, and novel classes. $mAP_{.50:.05:.95}$ is computed using mask labels for the COCO_{half} + LVIS setting and using bounding box labels for the LVIS + VG setting due to the absence of mask labels in the VG dataset.

Method	mAP _{all}	mAP _{known}	mAP _{novel}
k -means	1.48	17.24	0.13
RNCDL w/ ran. init.	1.87	22.85	0.09
ORCA	3.19	21.61	1.63
UNO	3.42	22.34	1.86
SimGCD	4.06	23.91	2.47
RNCDL	6.69	25.21	5.16
μ GCD [†]	4.93	25.36	3.58
NCDLR [†]	9.42	27.81	8.06
Ours	12.85	35.57	11.24

Table 1: Comparison for the COCO_{half} + LVIS setting. [†] indicates that only the GCD model in our method is replaced by the corresponding GCD method while all other components of our method remain unchanged.

Method	mAP _{all}	mAP _{known}	mAP _{novel}
RNCDL w/ ran. init.	2.13	7.71	0.81
SimGCD	2.37	9.13	0.92
RNCDL	4.46	12.55	2.56
μ GCD [†]	2.64	9.74	0.98
NCDLR [†]	4.71	12.87	2.63
Ours	5.21	13.28	3.27

Table 2: Quantitative comparison for the LVIS + VG setting.

4.2 Result

Tables 1 and 2 present quantitative comparisons for the COCO_{half} + LVIS and LVIS + VG settings, respectively. We compare the proposed method with the k -means (MacQueen et al. 1967) baseline and previous works including ORCA (Cao, Brbic, and Leskovec 2022), UNO (Fini et al. 2021), SimGCD (Wen, Zhao, and Qi 2023), and RNCDL (Fomenko et al. 2022). Additionally, we report results obtained by replacing the GCD model in our framework with recent GCD methods: μ GCD (Vaze, Vedaldi, and Zisserman 2023) and NCDLR (Zhang, Xu, and He 2023). The results demonstrate that our method significantly outperforms the previous state-of-the-art method (Fomenko et al. 2022) as well as methods focusing on balanced and curated datasets (Cao, Brbic, and Leskovec 2022; Fini et al. 2021; Wen, Zhao, and Qi 2023), across all metrics and settings. “RNCDL w/ ran. init.” refers to the method with random initialization in (Fomenko et al. 2022).

Figure 4 shows qualitative comparisons with RNCDL (Fomenko et al. 2022) for both settings.

4.3 Analysis

We conducted all the ablation studies and analyses using the COCO_{half} + LVIS setting unless otherwise noted. Table 3 presents an ablation study of the proposed components. The baseline model is constructed by using a fixed value for the temperature parameters for all samples in Eq. (4), excluding SAM, and treating all pseudo-labels as equivalent to human annotations. The results demonstrate that each module provides additional improvements across all metrics.

Table 4 compares the proposed ITA method with the TS

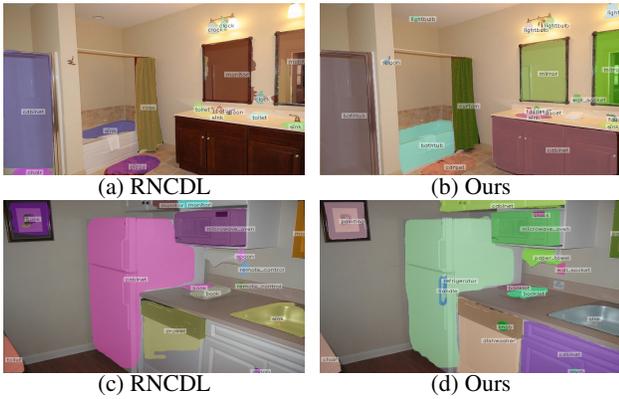


Figure 4: Qualitative results. (a) and (b) are $\text{COCO}_{half} + \text{LVIS}$ setting; (c) and (d) are $\text{LVIS} + \text{VG}$ setting.

Method	mAP _{all}	mAP _{known}	mAP _{novel}
Baseline	5.79	24.96	4.19
+ ITA	7.64	26.84	6.09
+ ITA + SAM	10.71	31.62	9.08
+ ITA + SAM + RDL	12.85	35.57	11.24

Table 3: Ablation study on the proposed components.

τ^{min}, τ^{max}	Method	mAP _{all}	mAP _{known}	mAP _{novel}
0.07, 0.5	TS	10.86	33.72	9.19
	Ours (ITA)	11.67	34.63	10.01
0.07, 1	TS	11.43	34.25	9.79
	Ours (ITA)	12.85	35.57	11.24

Table 4: Comparison on temperature assignment methods.

method from (Kukleva et al. 2023). Although the TS method is also designed for imbalanced data, our ITA method consistently outperforms it for both the selected hyperparameters (0.07, 1) and an alternative set (0.07, 0.5). Additionally, the table shows that the chosen hyperparameters yield better performance than the alternative set.

Table 5 compares the proposed SAM with previous attention methods, including MGCAM (Song et al. 2018), MGFP (Wang, Yu, and Gao 2021), CBAM-S (Woo et al. 2018), and RGA-S (Zhang et al. 2020). The results for the other methods are obtained by replacing only the SAM in the proposed method with the corresponding methods. The results demonstrate that our SAM outperforms all other methods across all metrics.

Table 6 compares our RDL method with the fixed and global reliability criteria in ST++ (Yang et al. 2020). The results confirm the importance of using selected pseudo-labels based on adjusted reliability criteria for each class throughout the training process.

Table 7 presents the results of our method using various class-agnostic instance mask generation models, including Mask R-CNN (He et al. 2017), OLN (Kim et al. 2022), LDET (Saito et al. 2022), UDOS (Kalluri et al. 2024), and

Method	mAP _{all}	mAP _{known}	mAP _{novel}
MGCAM	10.34	29.37	8.79
MGFPM	10.63	31.15	9.01
CBAM-S	11.19	31.83	9.56
RGA-S	11.62	33.48	9.98
Ours (SAM)	12.85	35.57	11.24

Table 5: Comparison with other attention methods.

Method	mAP _{all}	mAP _{known}	mAP _{novel}
$r = 25\%$	11.56	32.45	9.92
$r = 50\%$	11.79	33.36	10.17
$r = 75\%$	11.63	32.79	10.01
Ours (RDL)	12.85	35.57	11.24

Table 6: Comparison of reliability-based dynamic learning with fixed criteria.

Method	mAP _{all}	mAP _{known}	mAP _{novel}
Mask R-CNN	11.82	32.74	10.13
OLN	12.07	33.56	10.52
LDET	12.36	34.63	10.74
UDOS	12.54	35.12	11.05
GGN	12.85	35.57	11.24

Table 7: Comparison of class-agnostic instance mask generation methods.

GGN (Wang et al. 2022). The results demonstrate that our method consistently outperforms the previous state-of-the-art method (Fomenko et al. 2022), regardless of the choice of mask generation model.

5 Conclusion

Towards open-world instance segmentation, we present a novel GCD method in instance segmentation. To address the imbalanced distribution of instances, we introduce the instance-wise temperature assignment method as well as class-wise and dynamic reliability criteria. The former aims to improve the embedding space for class discovery, and the criteria are designed to effectively utilize pseudo-labels from the GCD model. Additionally, we propose an efficient soft attention module. The experimental results in two settings demonstrate that the proposed method outperforms previous methods by effectively discovering novel classes and segmenting instances of both known and novel categories.

Regarding limitations, this work assumes the full availability of labeled and unlabeled datasets from the beginning. Thus, it is suboptimal for scenarios where data is provided sequentially, such as in robot navigation. Additionally, we assume prior knowledge of the total number of classes, following most previous works.

Acknowledgments

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2023-00252434).

References

- An, W.; Tian, F.; Zheng, Q.; Ding, W.; Wang, Q.; and Chen, P. 2023. Generalized Category Discovery with Decoupled Prototypical Network. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11): 12527–12535.
- Arpit, D.; Jastrzebski, S.; Ballas, N.; Krueger, D.; Bengio, E.; Kanwal, M. S.; Maharaj, T.; Fischer, A.; Courville, A.; Bengio, Y.; and Lacoste-Julien, S. 2017. A Closer Look at Memorization in Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 233–242. PMLR.
- Bai, J.; Liu, Z.; Wang, H.; Chen, R.; Mu, L.; Li, X.; Zhou, J. T.; FENG, Y.; Wu, J.; and Hu, H. 2023. Towards Distribution-Agnostic Generalized Category Discovery. In *Advances in Neural Information Processing Systems*, volume 36, 58625–58647.
- Bellver, M.; Salvador, A.; Torres, J.; and Giro-i Nieto, X. 2019. Budget-aware Semi-Supervised Semantic and Instance Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Berrada, T.; Couprie, C.; Alahari, K.; and Verbeek, J. 2024. Guided Distillation for Semi-Supervised Instance Segmentation. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 464–472.
- Cao, K.; Brbic, M.; and Leskovec, J. 2022. Open-World Semi-Supervised Learning. In *International Conference on Learning Representations*.
- Caron, M.; Touvron, H.; Misra, I.; Jegou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9630–9640.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshick, R. 2022. Masked-attention Mask Transformer for Universal Image Segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1280–1289.
- Cuturi, M. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26.
- Fini, E.; Sangineto, E.; Lathuilière, S.; Zhong, Z.; Nabi, M.; and Ricci, E. 2021. A Unified Objective for Novel Class Discovery. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 9264–9272.
- Fomenko, V.; Elezi, I.; Ramanan, D.; Leal-Taixé, L.; and Osep, A. 2022. Learning to Discover and Detect Objects. In *Advances in Neural Information Processing Systems*, volume 35, 8746–8759.
- Gupta, A.; Dollár, P.; and Girshick, R. 2019. LVIS: A Dataset for Large Vocabulary Instance Segmentation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5351–5359.
- Han, K.; Vedaldi, A.; and Zisserman, A. 2019. Learning to Discover Novel Visual Categories via Deep Transfer Clustering. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 8400–8408.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9726–9735.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask R-CNN. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2980–2988.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Jain, J.; Li, J.; Chiu, M.; Hassani, A.; Orlov, N.; and Shi, H. 2023. OneFormer: One Transformer to Rule Universal Image Segmentation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2989–2998.
- Kalluri, T.; Wang, W.; Wang, H.; Chandraker, M.; Torresani, L.; and Tran, D. 2024. Open-world Instance Segmentation: Top-down Learning with Bottom-up Supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2693–2703.
- Kim, D.; Lin, T.-Y.; Angelova, A.; Kweon, I. S.; and Kuo, W. 2022. Learning Open-World Object Proposals Without Learning to Classify. *IEEE Robotics and Automation Letters*, 7(2): 5453–5460.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; Bernstein, M. S.; and Fei-Fei, L. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision*, 123(1): 32–73.
- Kuhn, H. W. 1955. The Hungarian Method for the Assignment Problem. *Naval Research Logistics Quarterly*, 2(1-2): 83–97.
- Kukleva, A.; Böhle, M.; Schiele, B.; Kuehne, H.; and Rupprecht, C. 2023. Temperature Schedules for Self-Supervised Contrastive Methods on Long-Tail Data. In *ICLR*.
- Lan, S.; Yu, Z.; Choy, C.; Radhakrishnan, S.; Liu, G.; Zhu, Y.; Davis, L. S.; and Anandkumar, A. 2021. DiscoBox: Weakly Supervised Instance Segmentation and Semantic Correspondence from Box Supervision. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3386–3396.
- Li, B.; Yao, Y.; Tan, J.; Zhang, G.; Yu, F.; Lu, J.; and Luo, Y. 2022. Equalized Focal Loss for Dense Long-Tailed Object Detection. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6980–6989.
- Li, Z.; Dai, B.; Simsek, F.; Meinel, C.; and Yang, H. 2023. ImbaGCD: Imbalanced Generalized Category Discovery. arXiv:2401.05353.

- Li, Z.; Meinel, C.; and Yang, H. 2023. Generalized Categories Discovery for Long-tailed Recognition. arXiv:2401.05352.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2999–3007.
- MacQueen, J.; et al. 1967. Some Methods for Classification and Analysis of Multivariate Observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, 281–297. Oakland, CA, USA.
- Pu, N.; Zhong, Z.; and Sebe, N. 2023. Dynamic Conceptual Contrastive Learning for Generalized Category Discovery. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7579–7588.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems*, volume 28.
- Riz, L.; Saltori, C.; Ricci, E.; and Poiesi, F. 2023. Novel Class Discovery for 3D Point Cloud Semantic Segmentation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9393–9402.
- Rosvall, M.; and Bergstrom, C. T. 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4): 1118–1123.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3): 211–252.
- Saito, K.; Hu, P.; Darrell, T.; and Saenko, K. 2022. Learning to Detect Every Thing in an Open World. In *Computer Vision – ECCV 2022*, 268–284. Cham: Springer Nature Switzerland. ISBN 978-3-031-20053-3.
- Song, C.; Huang, Y.; Ouyang, W.; and Wang, L. 2018. Mask-Guided Contrastive Attention Model for Person Re-identification. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1179–1188.
- Vaze, S.; Han, K.; Vedaldi, A.; and Zisserman, A. 2022. Generalized Category Discovery. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7482–7491.
- Vaze, S.; Vedaldi, A.; and Zisserman, A. 2023. No Representation Rules Them All in Category Discovery. In *Advances in Neural Information Processing Systems*, volume 36, 19962–19989.
- Wang, J.; Yu, X.; and Gao, Y. 2021. Mask Guided Attention For Fine-Grained Patchy Image Classification. In *2021 IEEE International Conference on Image Processing (ICIP)*, 1044–1048.
- Wang, W.; Feiszli, M.; Wang, H.; Malik, J.; and Tran, D. 2022. Open-World Instance Segmentation: Exploiting Pseudo Ground Truth From Learned Pairwise Affinity. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4412–4422.
- Wang, X.; Zhang, R.; Kong, T.; Li, L.; and Shen, C. 2020. SOLOv2: Dynamic and Fast Instance Segmentation. In *Advances in Neural Information Processing Systems*, volume 33, 17721–17732.
- Wang, Z.; Li, Y.; and Wang, S. 2022. Noisy Boundaries: Lemon or Lemonade for Semi-supervised Instance Segmentation? In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16805–16814.
- Wen, X.; Zhao, B.; and Qi, X. 2023. Parametric Classification for Generalized Category Discovery: A Baseline Study. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 16544–16554.
- Weng, Z.; Ogut, M. G.; Limonchik, S.; and Yeung, S. 2021. Unsupervised Discovery of the Long-Tail in Instance Segmentation Using Hierarchical Self-Supervision. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2603–2612.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional Block Attention Module. In *Computer Vision – ECCV 2018*, 3–19. Cham: Springer International Publishing. ISBN 978-3-030-01234-2.
- Yang, F.; Sun, Q.; Jin, H.; and Zhou, Z. 2020. Super-pixel Segmentation With Fully Convolutional Networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13961–13970.
- Yang, L.; Li, H.; Wu, Q.; Meng, F.; Qiu, H.; and Xu, L. 2023. Bias-Correction Feature Learner for Semi-Supervised Instance Segmentation. *IEEE Transactions on Multimedia*, 25: 5852–5863.
- Yang, L.; Zhuo, W.; Qi, L.; Shi, Y.; and Gao, Y. 2022. ST++: Make Self-training Work Better for Semi-supervised Semantic Segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4258–4267.
- Zhang, C.; Xu, R.; and He, X. 2023. Novel Class Discovery for Long-tailed Recognition. *Transactions on Machine Learning Research*.
- Zhang, D.; Nan, F.; Wei, X.; Li, S.-W.; Zhu, H.; McKeown, K.; Nallapati, R.; Arnold, A. O.; and Xiang, B. 2021. Supporting Clustering with Contrastive Learning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 5419–5430. Online: Association for Computational Linguistics.
- Zhang, S.; Khan, S.; Shen, Z.; Naseer, M.; Chen, G.; and Khan, F. S. 2023. PromptCAL: Contrastive Affinity Learning via Auxiliary Prompts for Generalized Novel Category Discovery. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3479–3488.
- Zhang, Z.; Lan, C.; Zeng, W.; Jin, X.; and Chen, Z. 2020. Relation-Aware Global Attention for Person Re-identification. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3183–3192.
- Zhao, Y.; Zhong, Z.; Sebe, N.; and Lee, G. H. 2022. Novel Class Discovery in Semantic Segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4330–4339.