

ImbaGCD: Imbalanced Generalized Category Discovery

Ziyun Li
Hasso Plattner Institute
University of Potsdam
ziyun.li@hpi.de

Ben Dai
Chinese University of HongKong
bendai@cuhk.edu.hk

Furkan Simsek
Hasso Plattner Institute
University of Potsdam
furkan.simsek@hpi.de

Christoph Meinel
Hasso Plattner Institute
University of Potsdam
christoph.meinel@hpi.de

Haojin Yang
Hasso Plattner Institute
University of Potsdam
haojin.yang@hpi.de

Abstract

Generalized class discovery (GCD) aims to infer known and unknown categories in an unlabeled dataset leveraging prior knowledge of a labeled set comprising known classes. Existing research implicitly/explicitly assumes that the frequency of occurrence for each category, whether known or unknown, is approximately the same in the unlabeled data. However, in nature, we are more likely to encounter known/common classes than unknown/uncommon ones, according to the long-tailed property of visual classes. Therefore, we present a challenging and practical problem, *Imbalanced Generalized Category Discovery (ImbaGCD)*, where the distribution of unlabeled data is imbalanced, with known classes being more frequent than unknown ones. To address these issues, we propose *ImbaGCD*, A novel optimal transport-based expectation maximization framework that accomplishes generalized category discovery by aligning the marginal class prior distribution. *ImbaGCD* also incorporates a systematic mechanism for estimating the imbalanced class prior distribution under the GCD setup. Our comprehensive experiments reveal that *ImbaGCD* surpasses previous state-of-the-art GCD methods by achieving an improvement of approximately 2 - 4% on CIFAR-100 and 15 - 19% on ImageNet-100, indicating its superior effectiveness in solving the Imbalanced GCD problem.

1. Introduction

Existing machine learning models can attain excellent performance when trained on large-scale datasets with human annotations. However, the success of these models is strongly dependent on the fact that they are only required to recognize images from the same set of classes with extensive human annotations on which they are train.

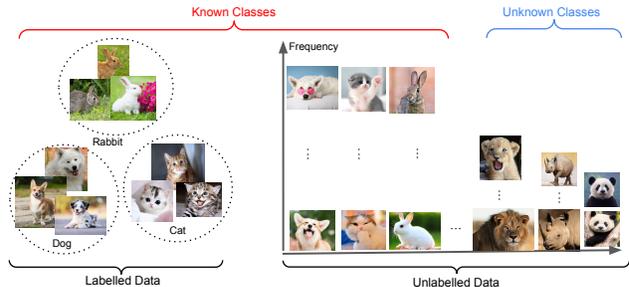


Figure 1. Illustration of ImbaGCD. ImbaGCD attempts to identify known categories and discover new classes within a vast amount of unlabeled data from the real world. In the real world, unlabeled data includes both known class from the labeled set and unknown class, where known classes (e.g., cat, dog, rabbit) dominate the well-represented “head,” while “unknown” classes (e.g., lion, rhino, panda) primarily reside in the underrepresented “tail” of the distribution.

This constraint restricts the applicability of these models in the real world where unannotated data from unseen categories may be encountered. To overcome this limitation, researchers have developed related research topics, such as semi-supervised learning [7, 16] that utilizes both labeled and unlabeled data to train a robust model, few-shot learning [40] that aims to generalize to new classes with limited annotated samples, open-set recognition [38] that identifies whether an unlabeled image belongs to one of the known classes, and novel category discovery (NCD) [14, 17, 18, 32] that partitions unannotated data from unknown categories by transferring knowledge from known ones. NCD initially assumed all unannotated images were from unknown categories, which is not realistic. GCD [44] is introduced to consider unannotated images from both known and un-

known categories, better reflecting real-world scenarios.

The Generalized Category Discovery (GCD) problem presents challenges due to the agnostic nature of real-world unlabeled data. Two primary concerns arise: (i) *limited knowledge of new categories’ appearances*, which complicates selecting labeled datasets with high semantic similarity. As indicated in [10, 31, 32], GCD may become ill-defined if known and novel classes do not share high-level semantic features. Furthermore, [31, 32] highlights that utilizing supervised knowledge from labeled sets can result in suboptimal performance in low semantic similarity situations. (ii) *Lack of information about the occurrence frequency* for each category, whether known or unknown, which constitutes the central focus of our research. However, existing research [4, 37, 41, 44] often overlooks the issue of category occurrence frequency and tends to use experimental setups based on [44], where the unknown class occurs twice as often as the known class. This setting, however, misrepresents real-world scenarios. In line with the long-tailed property of visual classes, we are more likely to encounter known classes in natural environments. As shown in Figure 1, known classes (e.g. cat, dog and rabbit) with their ease of label acquisition, dominate the well-represented “head” of the distribution, while “unknown” classes (e.g., lion, rhino and panda) are harder to obtain, primarily residing in the underrepresented “tail”.

Therefore, we present Imbalanced Generalized Category Discovery (ImbaGCD), a realistic problem where the distribution of unlabeled data is imbalanced, with known classes more frequent than unknown ones. Our work focuses on two primary challenges: (i) *Estimating indeterminate class prior*. Unlike conventional long-tailed learning (LTL) scenarios [2, 5, 33], most existing methods [8, 11, 19, 21, 22] rely on supervised learning and labeled data for handling class imbalance. This renders the estimation of the marginal class prior distribution particularly challenging, as it cannot be achieved by simply counting training samples based on class labels. In this work, we address this by an iterative class prior estimation technique, which serves as a robust approximation of the true prior (Section 3.4). (ii) *Mitigating bias towards head classes*, which is a key challenge in imbalanced data settings [24], and is also a primary concern in ImbaGCD. Specifically, tail/unknown samples are often misclassified as head/known classes, as models tend to focus on learning patterns from prevalent classes, leading to poorer performance on minority classes. Our approach involves applying constraints to pseudo-labels to align their distribution with the class prior distribution, promoting the identification of unlabeled samples as unknown classes. By formulating this as an optimal transport problem [45], we efficiently solve the constrained optimization objective using the Sinkhorn-Knopp algorithm [12].

We carry out a comprehensive evaluation of ImbaGCD

on benchmark datasets, outperforming the state-of-the-art by margins of approximately 2-4% and 15-19% on CIFAR-100 and ImageNet-100, respectively, across multiple imbalanced settings. In addition, ImbaGCD exhibits competitive performance in both balanced and original GCD settings, underscoring its versatility and effectiveness in a variety of situations. Our contributions are summarized as follows:

- We present a challenging and practical problem, Imbalanced Generalized Category Discovery (ImbaGCD), where the distribution of unlabeled data is imbalanced, with known classes being more frequent than unknown ones.
- We propose ImbaGCD, a novel optimal transport-based expectation maximization framework that enables the discovery of generalized classes by matching the marginal class prior distribution. ImbaGCD also incorporates a systematic mechanism for estimating the imbalance class prior distribution under the GCD setup.
- Our extensive experiments demonstrate that ImbaGCD outperforms previous state-of-the-art GCD methods on standardized benchmarks, indicating its superior effectiveness in solving the ImbaGCD problem.

2. PRELIMINARIES

In this section, we first introduce the GCD setup, and then briefly review the optimal transport.

2.1. Problem Setup

We denote (\mathbf{X}_l, Y_l) and (\mathbf{X}_u, Y_u) as random samples under the *labeled/unlabeled probability measures* $\mathbb{P}_{\mathbf{X}, Y}$ and $\mathbb{Q}_{\mathbf{X}, Y}$, respectively. $\mathbf{X}_l \in \mathcal{X}_l \subset \mathbb{R}^d$ and $\mathbf{X}_u \in \mathcal{X}_u \subset \mathbb{R}^d$ are the labeled/unlabeled feature vectors, $Y_l \in \mathcal{Y}_l$ and $Y_u \in \mathcal{Y}_u$ are the true labels of labeled/unlabeled data, where \mathcal{Y}_l and \mathcal{Y}_u are the label sets under the labeled and unlabeled probability measures $\mathbb{P}_{\mathbf{X}, Y}$ and $\mathbb{Q}_{\mathbf{X}, Y}$, respectively.

Definition 1 (Generalized Category Discovery) *Let $\mathbb{P}_{\mathbf{X}_l, Y_l}$ be a labeled probability measure on $\mathcal{X}_l \times \mathcal{Y}_l$, and $\mathbb{Q}_{\mathbf{X}_u, Y_u}$ be an unlabeled probability measure on $\mathcal{X}_u \times \mathcal{Y}_u$, with $\mathcal{Y}_l \subset \mathcal{Y}_u$. (\mathcal{Y}_u comprises both known classes \mathcal{Y}_l and unknown classes $\mathcal{Y}_n = \mathcal{Y}_u \setminus \mathcal{Y}_l$). Given a labeled dataset \mathcal{L}_n sampled from $\mathbb{P}_{\mathbf{X}_l, Y_l}$ and an unlabeled dataset \mathcal{U}_n sampled from $\mathbb{Q}_{\mathbf{X}_u}$, GCD aims to predict the label Y_u for each unlabeled instance X_u , which may belong to either known or unknown classes.*

Specifically, under our proposed imbalanced GCD, as discussed before, the class distributions in the unlabeled set are skewed with $\mathbb{P}_{Y_l}(i) > \mathbb{P}_{Y_n}(j)$ for all $i \in \mathcal{Y}_l$ and $j \in \mathcal{Y}_n$. This problem formulation acknowledges the class distribution differences between known and unknown categories in

unlabeled sets, instead of assuming $\mathbb{P}_{Y_l}(i) \approx \mathbb{P}_{Y_n}(j)$ for all $i \in \mathcal{Y}_l$ and $j \in \mathcal{Y}_n$.

2.2. Reminders on Optimal Transport

Optimal transport (OT) is a method to quantify the cost of converting one probability measure to another, which offers a unique perspective to understand imbalanced problems. Given two random variables \mathbf{X} and Y , their corresponding probability measures are denoted as \mathbf{r} and \mathbf{w} . Furthermore, the cost function $\mathcal{C}(\mathbf{X}, Y) : \mathbf{X} \times Y \rightarrow \mathbb{R}_+$ represents the expense of transferring \mathbf{X} to Y . Consequently, the OT distance between \mathbf{X} and Y can be defined as follows:

$$\text{OT}(\mathbf{r}, \mathbf{w}) = \min_{\pi \in \Pi(\mathbf{r}, \mathbf{w})} \int_{\mathbf{X} \times Y} \mathcal{C}(\mathbf{x}, y) \pi(\mathbf{x}, y) d\mathbf{x} dy$$

$$\Pi(\mathbf{r}, \mathbf{w}) := \left\{ \int_Y \pi(\mathbf{x}, y) dy = \mathbf{r}(\mathbf{x}), \int_X \pi(\mathbf{x}, y) d\mathbf{x} = \mathbf{w}(y) \right\}$$

where $\pi(\mathbf{r}, \mathbf{w})$ is the joint probability measure with \mathbf{r} and \mathbf{w} [46]. In empirical version, the OT distance can be expressed using discrete distributions and a cost matrix \mathbf{M} :

$$d_M(\mathbf{r}, \mathbf{w}) = \min_{\mathbf{P} \in U(\mathbf{r}, \mathbf{w})} \langle \mathbf{P}, \mathbf{M} \rangle$$

$$U(\mathbf{r}, \mathbf{w}) := \{ \mathbf{P} \in \mathbb{R}_+^{d \times d} \mid \mathbf{P} \mathbf{1}_d = \mathbf{r}, \mathbf{P}^T \mathbf{1}_d = \mathbf{w} \}$$

where $U(\mathbf{r}, \mathbf{w})$ represents the transport polytope of \mathbf{r} and \mathbf{w} , namely the polyhedral set of $d \times d$ matrices with non-negative entries whose rows and columns sum to \mathbf{r} and \mathbf{w} , respectively. The primary objective of OT is to determine a transportation matrix \mathbf{P} that minimizes the distance $d_M(\mathbf{r}, \mathbf{w})$. Although OT serves as a distance measure between probability distributions under a specific cost matrix, solving the optimization problem with network simplex or interior point methods can be computationally demanding. To address this challenge, OT with entropy constraint has been introduced, which optimizes with a lower computational cost while maintaining sufficient smoothness [12]. By incorporating a Lagrangian multiplier for the entropy constraint, the new formulation is defined as:

$$d_M^\lambda(\mathbf{r}, \mathbf{w}) = \langle \mathbf{P}^\lambda, \mathbf{M} \rangle$$

where $\mathbf{P}^\lambda = \arg \min_{\mathbf{P} \in U(\mathbf{r}, \mathbf{w})} \langle \mathbf{P}, \mathbf{M} \rangle - \lambda h(\mathbf{P})$,

$\lambda > 0$, $h(\mathbf{P}) = -\sum_{n=1}^N \sum_{k=1}^K \mathbf{P}_{nk} \log \mathbf{P}_{nk}$, and $d_M^\lambda(\mathbf{r}, \mathbf{w})$ is also referred to as dual-Sinkhorn divergence. The matrix scaling algorithms can compute this divergence with reduced computational demand. A lemma ensures the convergence and uniqueness of the solution.

Lemma 1 For $\lambda > 0$, the solution \mathbf{P}^λ is unique and can be represented as $\mathbf{P}^\lambda = \text{diag}(\alpha) \mathbf{K} \text{diag}(\beta)$, where α and β are two non-negative vectors uniquely determined up to a multiplicative factor, and $\mathbf{K} = e^{-\mathbf{M}/\lambda}$ is the element-wise exponential of $-\mathbf{M}/\lambda$.

The lemma above establishes the uniqueness of the solution \mathbf{P}^λ [39]. Additionally, \mathbf{P}^λ can be efficiently computed using Sinkhorn's fixed point iteration: $\alpha, \beta \leftarrow \mathbf{r} ./ \mathbf{K} \beta, \mathbf{w} ./ \mathbf{K}^T \alpha$, where $./$ denotes element-wise division.

Algorithm 1: Pseudo-code for ImbaGCD

Input: Training dataset D , classifier f , uniform marginal \mathbf{r}, \mathbf{w} , and hyperparameters $\lambda_{proto}, \lambda_{sup}, \tau, \mu$

- 1 **Algorithm** ImbaGCD ($\lambda_{proto}, \lambda_{sup}, \tau, \mu$):
- 2 **for** $epoch = 1, 2, \dots$, **do**
- 3 **for** $step = 1, 2, \dots$, **do**
- 4 //E-step: estimate pseudo-label matrix \mathbf{A} :
- 5 Get classifier prediction \mathbf{P} on a mini-batch of the unlabeled data of size B_u
- 6 Calculate \mathbf{K} such that $k_{ij} = p_{ij}^\lambda$
- 7 **for** $t=1, \dots, T$ **do**
- 8 // Sinkhorn's fixed point iteration
- 9 $\alpha \leftarrow \mathbf{w} ./ (\mathbf{K} \beta), \quad \beta \leftarrow \mathbf{r} ./ (\mathbf{K}^T \alpha)$
- 10 **end**
- 11 $\mathbf{A} = B_u \text{diag}(\alpha) \mathbf{K} \text{diag}(\beta)$
- 12 // M-step: update model parameters:
- 13 $\theta^k = \text{SGD}(\mathcal{L}_{\text{overall}}, \theta^{k-1})$
- 14 // Overall loss function is provided Eq. 11
- 15 **end**
- 16 $\mathbf{r} \leftarrow \mu \mathbf{r} + (1 - \mu) \mathbf{z} / \text{update the class prior}$
- 17 $\mathbf{c}_k \leftarrow \mu \mathbf{c}_k + (1 - \mu) \mathbf{v}_k / \text{update the prototype}$
- 18 **end**

3. Method

Our proposed method, ImbaGCD, is based on the Expectation-Maximization (EM) algorithm, which assigns pseudo-labels using the Sinkhorn algorithm in the E-step and updates the model with contrastive learning in the M-step. We iteratively update the class prior and prototype every epoch. The detailed process is outlined in Algorithm 1.

3.1. Prototype Loss for Unlabeled Data

Our objective is to find the network parameters θ that maximizes the log-likelihood function of the observed m unlabeled samples: $\theta^* = \arg \max_{\theta} \sum_{i=1}^m \log p(\mathbf{x}^{(i)} \mid \theta)$.

We assume that the observed data $\{\mathbf{x}^{(i)}\}_{i=1}^m$ are related to latent variable $\mathbf{C} = \{\mathbf{c}_k\}_{k=1}^{|\mathcal{Y}_u|}$ which denotes the prototypes of the data, where $|\mathcal{Y}_u|$ denotes the total number of unknown classes ($\mathcal{Y}_l \subset \mathcal{Y}_u$). The joint distribution of $\mathbf{x}^{(i)}$ and $\mathbf{c}^{(i)}$ is given as follow:

$$\theta^* = \arg \max_{\theta} \sum_{i=1}^m \log \sum_{\mathbf{c}_k \in \mathcal{C}} p(\mathbf{x}^{(i)}, \mathbf{c}_k^{(i)} \mid \theta) \quad (1)$$

It is hard to optimize this function directly, so we use a surrogate function to lower-bound it:

$$\begin{aligned}
& \sum_{i=1}^m \log \sum_{c_k \in C} p(x^{(i)}, c_k^{(i)} | \theta) \\
&= \sum_{i=1}^m \log \sum_{c_k \in C} Q(c_k^{(i)}) \frac{p(x^{(i)}, c_k^{(i)} | \theta)}{Q(c_k^{(i)})} \quad (2) \\
&\geq \sum_{i=1}^m \sum_{c_k \in C} Q(c_k^{(i)}) \log \frac{p(x^{(i)}, c_k^{(i)} | \theta)}{Q(c_k^{(i)})}
\end{aligned}$$

To achieve equality, we have $Q(c_k^{(i)}) = p(c_k^{(i)} | \mathbf{x}^{(i)}, \theta)$. By ignoring the constant, we aim to maximize:

$$\sum_{i=1}^m \sum_{c_k \in C} Q(c_k^{(i)}) \log p(\mathbf{x}^{(i)}, c_k^{(i)} | \theta) \quad (3)$$

E-step The aim of this step is to estimate the value of $Q(c_k^{(i)})$, which can be represented as $p(c_k^{(i)} | \mathbf{x}^{(i)}, \theta)$. To this end, the Sinkhorn algorithm [12] is employed to generate pseudo labels in line with the prior class distribution, as opposed to maximizing the prediction with $\hat{y}^{(i)} = \operatorname{argmax}_{c_k \in C} \mathbf{c}_k^\top \cdot f_\theta(\mathbf{x}^{(i)})$. In order to formalize the optimal transport problem for proper label assignments, consider the following setup. At each training step, we aim to search for pseudo-labels \mathbf{A} that closely approximate the current classifier’s predictions \mathbf{P} , while adhering to specific constraints:

$$\begin{aligned}
& \min_{\mathbf{A} \in \Delta} E(\mathbf{A}, \mathbf{P}) = \langle \mathbf{A}, -\log(\mathbf{P}) \rangle \\
& \text{s.t. } \Delta = \{ \mathbf{A}^\top \mathbf{1}_m = \mathbf{r}, \mathbf{A} \mathbf{1}_L = \mathbf{w} \} \quad (4)
\end{aligned}$$

where \mathbf{r} is an $|\mathcal{Y}_u|$ -dimensional probability simplex that indicates the prior class distribution. Note that, here we temporarily assume we have a decent estimation of the class priors and we will describe the means of estimation in Section 3.4. The column vector $\mathbf{w} = \frac{1}{m} \mathbf{1}_m$ indicates that our m training examples are sampled uniformly. To resolve Eq. 4, we adapt the well-known Sinkhorn-Knopp algorithm [12] for efficient optimization. Formally, we define a matrix \mathbf{K} such that $\mathbf{K}_{ij} = \mathbf{P}_{ij}^\lambda$, where $\lambda > 0$ is a smoothing regularization coefficient. \mathbf{K} can be efficiently computed using Sinkhorn’s fixed point iteration: $\alpha, \beta \leftarrow \mathbf{r} / \mathbf{K} \beta, \mathbf{w} / \mathbf{K}^\top \alpha$, where $\cdot /$ denotes element-wise division. Additionally, inspired by [6, 20], we also involve a queue acceleration trick to avoid traversing the whole training set.

M-step Based on the E-step, we are ready to maximize the lower-bound in Eq. 3 with respect to θ :

$$\begin{aligned}
& \sum_{i=1}^n \sum_{c_k \in C} Q(c_k^{(i)}) \log p(\mathbf{x}^{(i)}, c_k^{(i)} | \theta) \\
&= \sum_{i=1}^n \sum_{c_k \in C} \underbrace{\mathbb{1}(\mathbf{x}^{(i)} \in c_k^{(i)})}_{\text{Achieved by Sinkhorn algorithm in E-step}} \log p(\mathbf{x}^{(i)}, c_k^{(i)} | \theta) \quad (5)
\end{aligned}$$

And we also have $p(\mathbf{x}^{(i)}, c_k^{(i)} | \theta) = p(\mathbf{x}^{(i)} | c_k^{(i)}, \theta) p(c_k^{(i)} | \theta)$, we derive prior probability $p(c_k^{(i)} | \theta)$ from class prior estimation (Section 3.4). Following [29], we assume an isotropic Gaussian distribution around each prototype with an the same variance σ :

$$\begin{aligned}
& p(x^{(i)} | c_k^{(i)}, \theta) \\
&= \exp\left(\frac{-(v_i - c_k^{(i)})^2}{2\sigma^2}\right) / \sum_{j=1}^{|\mathcal{Y}_u|} \exp\left(\frac{-(v_i - c_j)^2}{2\sigma^2}\right), \quad (6)
\end{aligned}$$

with $\mathbf{v}^{(i)} = f_\theta(\mathbf{x}^{(i)})$ and $\mathbf{x}^{(i)} \in c_k^{(i)}$, combining the above equations, we express maximum log-likelihood estimation as:

$$\theta^* = \operatorname{argmin}_\theta \sum_{i=1}^m -\log \frac{\exp(\mathbf{v}^{(i)} \cdot \mathbf{c}_k^{(i)})}{\sum_{j=1}^{|\mathcal{Y}_u|} \exp(\mathbf{v}^{(i)} \cdot \mathbf{c}_j)} - \log p(c_k^{(i)} | \theta) \quad (7)$$

The prototype loss of the unlabeled data is:

$$\mathcal{L}_{proto}^{(i)} = -\log \frac{\exp(\mathbf{v}^{(i)} \cdot \mathbf{c}_k^{(i)})}{\sum_{j=1}^{|\mathcal{Y}_u|} \exp(\mathbf{v}^{(i)} \cdot \mathbf{c}_j)} - \log p(c_k^{(i)} | \theta) \quad (8)$$

3.2. Representation Improvement

To enhance model representation, we adopt unsupervised contrastive learning [9, 20] for unlabeled data and supervised contrastive learning [25] for labeled data, as in [4, 41]. Specifically, let \mathbf{v}_i and \mathbf{v}'_i represent features from two views (random augmentations) of the same image within a mini-batch B . The instance-level unsupervised contrastive loss and the supervised contrastive loss are defined as follows:

$$\begin{aligned}
\mathcal{L}_{ins}^{(i)} &= \frac{1}{|B_u|} \sum_{i \in B_u} -\log \frac{\exp(\mathbf{v}_i \cdot \mathbf{v}'_i / \tau)}{\sum_{i \neq j} \exp(\mathbf{v}_i \cdot \mathbf{v}_j / \tau)}, \quad (9) \\
\mathcal{L}_{sup}^{(i)} &= \frac{1}{|B_l|} \sum_{i \in B_l} \frac{1}{|\mathcal{N}^{(i)}|} \sum_{q \in \mathcal{N}^{(i)}} -\log \frac{\exp(\mathbf{v}_i \cdot \mathbf{v}_q / \tau)}{\sum_{i \neq j} \exp(\mathbf{v}_i \cdot \mathbf{v}_j / \tau)}, \quad (10)
\end{aligned}$$

where B_u and B_l represent the unlabeled and labeled subsets of mini-batch B , respectively. Furthermore, $\mathcal{N}^{(i)}$ refers to the indices of other images in the batch sharing the same label, and τ signifies a temperature parameter.

3.3. Overall Loss Objective

The overall loss objective is a weighted sum of the unsupervised and supervised contrastive losses:

$$\mathcal{L}_{\text{overall}} = \underbrace{\mathcal{L}_{\text{ins}} + \lambda_{\text{proto}} \mathcal{L}_{\text{proto}}}_{\mathcal{L}_{\text{unlabeled}}} + \underbrace{\lambda_{\text{sup}} \mathcal{L}_{\text{sup}}}_{\mathcal{L}_{\text{labeled}}}, \quad (11)$$

where λ_{proto} and λ_p are a hyper-parameter controlling the relative weight of the supervised loss. The contrastive loss derived from the unlabeled data can be examined within two distinct hierarchical tiers: the instance level, referred to as instance contrastive loss, \mathcal{L}_{ins} , and the category level, identified as prototype contrastive loss, $\mathcal{L}_{\text{proto}}$. Conversely, for the labeled data, only the category level is taken into account, which is referred to as supervised contrastive loss, \mathcal{L}_{sup} .

3.4. Class Prior Estimation & Prototype Updates

Moving-average distribution update We suggest employing model predictions for class prior estimation following [15]. However, due to potential inaccuracies and biases in early training stages, we propose a moving-average update mechanism to enhance reliability. Starting with a uniform class prior $\mathbf{r} = [1/C, \dots, 1/C]$, we iteratively refine the distribution per epoch.

$$\mathbf{r} := \mu \mathbf{r} + (1 - \mu) \mathbf{z},$$

$$\text{where } \mathbf{z}_j = \frac{1}{n} \sum_{i=1}^n \mathbb{1} \left(j = \arg \max_{j'} f_{j'}(\mathbf{x}_i) \right), \mu \in [0, 1]$$

The class prior is continuously updated through a linear function, resulting in more stable training dynamics. As the training progresses, the model’s accuracy improves, making the estimated distribution increasingly dependable.

Momentum prototypes update A canonical approach for updating prototype embeddings is computationally expensive. To reduce training latency, we employ a moving-average strategy [28] for updating class-conditional prototype vectors:

$$\mathbf{c}_k := \mu \mathbf{c}_k + (1 - \mu) \mathbf{v}_k,$$

where the prototype \mathbf{c}_k of class k -th is defined by the moving average of the normalized embeddings \mathbf{v}_k , whose predicted class conforms to k . For both iterate updates, we employ the same hyperparameter μ and we update prototypes and class prior each epoch.

4. Experiments

In this section, we conduct an experimental analysis of the proposed ImbaGCD approach under both original and varying imbalanced scenarios.

4.1. Setup

Datasets We conducted experiments on three datasets, including CIFAR10, CIFAR100 and ImageNet-100 (where ImageNet-100 is a subsampled version of the ImageNet dataset with 100 classes), the classes were divided into 50% known and 50% unknown classes. And we randomly selected 50% of the known class samples as the labeled dataset, which is balanced. However, the sample sizes of known and unknown classes in the unlabeled set are imbalanced, which is adjusted based on the imbalanced factor ρ . Here, ρ was defined as the ratio of the sample sizes of the known and unknown classes, i.e., $\rho = \frac{n_k}{n_u}$, where n_k and n_u represent the sample sizes of known and unknown classes, respectively, and $\rho \in \{0.5, 1, 5, 10\}$. It is worth noting that previous works often selected 50% of the known class samples as the labeled dataset, and the remaining samples were used for the unlabeled set, resulting in a fixed imbalance ratio of $\rho = 0.5$.

Evaluation metrics We adopt the evaluation strategy outlined in [4, 41, 44] and report the following metrics: (1) overall accuracy across all classes, (2) classification accuracy for known classes. Additionally, we evaluate novel data using two distinct evaluation settings: (3) **unknown-aware** and (4) **unknown-agnostic**. As described in [41], the accuracy for novel classes and all classes is determined by solving an optimal assignment problem using the Hungarian algorithm [27]. In the unknown-aware evaluation, we adhere to the established GCD methods by first isolating all unlabeled samples associated with unknown classes. We then perform clustering specifically within these unknown class categories. However, this approach may not accurately reflect real-world scenarios, as directly distinguishing between known and unknown classes within unlabeled data is often impractical. To address this limitation, we also report unknown-agnostic accuracy, which refrains from using any information to differentiate between known and unknown classes within unlabeled data, providing a more realistic and unbiased evaluation metric.

Implementation details We utilize ResNet-18 as the backbone architecture for CIFAR-100 and ResNet-50 for ImageNet-100. In addition, we introduce a trainable two-layer MLP projection head that maps the features from the penultimate layer to a lower-dimensional space \mathbb{R}^d ($d = 128$). This projection technique has proven effective for contrastive loss [9]. In line with the methods proposed in [4, 41], we implement regularization by calculating the KL-divergence between the predicted label distribution and the class prior, which helps to improve model stability.

For both CIFAR10/100 and ImageNet-100, the model undergoes training for 80 and 120 epochs, respectively, with

Table 1. Performance comparison of various methods on CIFAR100 under different imbalanced factors ρ ($\rho = 0.5$ is original GCD setting). Our method consistently outperforms others on novel class, showcasing its effectiveness in handling class imbalance. We report the mean and standard deviation of the clustering accuracy across 3 runs for multiple methods. The higher mean value is presented in bold, while the results within standard deviation of the average accuracy are not bolded.

IMF	Metrics	Methods			
		GCD	ORCA	OpenCon	Ours
$\rho = 0.5$	All	45.41 \pm 0.13	55.68 \pm 0.33	51.85 \pm 0.63	53.51 \pm 0.26
	Known	67.61 \pm 0.12	66.41 \pm 0.31	69.07 \pm 0.29	68.09 \pm 0.13
	Unknown-aware	34.31 \pm 0.22	42.63 \pm 0.67	45.76 \pm 0.32	47.92 \pm 0.33 (+ 2.16)
	Unknown-agnostic	18.12 \pm 0.34	38.95 \pm 0.76	42.11 \pm 0.54	46.22 \pm 0.33 (+ 4.11)
$\rho = 1$	All	48.36 \pm 0.08	47.37 \pm 0.52	53.20 \pm 0.33	54.06 \pm 0.45
	Known	71.48 \pm 0.24	65.14 \pm 0.17	68.00 \pm 0.06	67.98 \pm 0.37
	Unknown-aware	25.24 \pm 0.07	34.93 \pm 1.04	43.58 \pm 0.34	43.39 \pm 0.59
	Unknown-agnostic	12.02 \pm 0.46	29.61 \pm 1.02	38.40 \pm 0.65	40.72 \pm 0.9 (+ 2.32)
$\rho = 5$	All	63.13 \pm 0.12	56.37 \pm 0.12	61.84 \pm 0.29	59.48 \pm 0.48
	Known	70.96 \pm 0.13	64.40 \pm 0.11	69.37 \pm 0.23	67.82 \pm 0.06
	Unknown-aware	24.04 \pm 0.12	25.36 \pm 0.40	35.06 \pm 0.34	37.87 \pm 1.59 (+ 2.21)
	Unknown-agnostic	7.18 \pm 0.34	16.18 \pm 0.23	24.21 \pm 0.67	27.64 \pm 2.26 (+ 3.43)
$\rho = 10$	All	66.21 \pm 0.23	59.61 \pm 0.27	65.05 \pm 0.29	63.21 \pm 0.05
	Known	70.36 \pm 0.31	64.27 \pm 0.32	69.80 \pm 0.23	67.82 \pm 0.07
	Unknown-aware	24.74 \pm 0.45	26.21 \pm 0.69	33.01 \pm 0.34	34.87 \pm 0.70 (+ 1.76)
	Unknown-agnostic	7.28 \pm 0.21	13.01 \pm 0.34	17.57 \pm 0.67	21.68 \pm 0.29 (+ 4.12)

Table 2. Performance comparison of various methods on ImageNet100 under different imbalanced factors ρ , ($\rho = 0.5$ is original GCD setting). Our method consistently outperforms others on novel class under imbalanced settings and achieves competitive performance on the balanced setting. We report the mean and standard deviation of the clustering accuracy across 3 runs for multiple methods. The higher mean value is presented in bold, while the results within standard deviation of the average accuracy are not bolded.

IMF	Metrics	Methods			
		GCD	ORCA	OpenCon	Ours
$\rho = 0.5$	All	77.12 \pm 0.56	74.93 \pm 0.34	82.22 \pm 0.56	81.90 \pm 0.67
	Known	87.02 \pm 0.36	89.21 \pm 0.06	90.65 \pm 0.04	91.19 \pm 0.15
	Unknown-aware	57.48 \pm 0.76	67.18 \pm 0.27	78.12 \pm 0.80	77.89 \pm 0.88
	Unknown-agnostic	42.68 \pm 0.77	65.43 \pm 0.37	78.01 \pm 0.83	77.79 \pm 0.97
$\rho = 1$	All	63.06 \pm 0.66	68.01 \pm 0.31	82.44 \pm 0.24	82.34 \pm 0.39
	Known	88.80 \pm 0.43	88.99 \pm 0.05	90.62 \pm 0.12	90.56 \pm 0.18
	Unknown-aware	37.29 \pm 1.11	47.68 \pm 0.70	74.45 \pm 0.36	74.56 \pm 0.79
	Unknown-agnostic	31.16 \pm 1.48	47.28 \pm 0.66	74.35 \pm 0.38	74.24 \pm 0.89
$\rho = 5$	All	77.29 \pm 0.17	76.79 \pm 0.13	81.87 \pm 0.23	83.01 \pm 0.11
	Known	87.51 \pm 0.21	89.02 \pm 0.12	90.77 \pm 0.10	88.89 \pm 0.06
	Unknown-aware	23.83 \pm 0.24	20.19 \pm 0.55	39.76 \pm 1.4	54.35 \pm 0.38 (+ 14.59)
	Unknown-agnostic	14.80 \pm 0.23	16.42 \pm 0.73	37.93 \pm 1.3	53.97 \pm 0.38 (+ 16.04)
$\rho = 10$	All	83.32 \pm 0.15	81.98 \pm 0.03	84.60 \pm 0.10	83.23 \pm 0.17
	Known	89.71 \pm 0.03	89.24 \pm 0.02	90.85 \pm 0.12	87.48 \pm 0.23
	Unknown-aware	22.17 \pm 0.18	16.94 \pm 0.07	26.70 \pm 0.11	42.62 \pm 0.50 (+ 15.92)
	Unknown-agnostic	11.47 \pm 0.28	10.31 \pm 0.22	22.89 \pm 0.27	41.23 \pm 0.58 (+ 18.34)

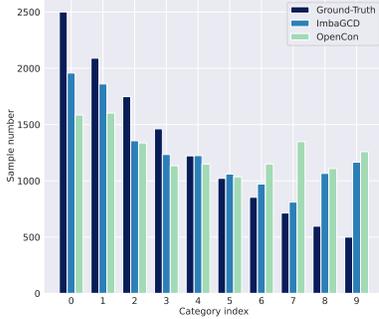


Figure 2. Comparison of the predicted sample numbers for each class on the CIFAR-10 dataset, using an exponential decreasing strategy with an imbalance factor (ρ) of 5.

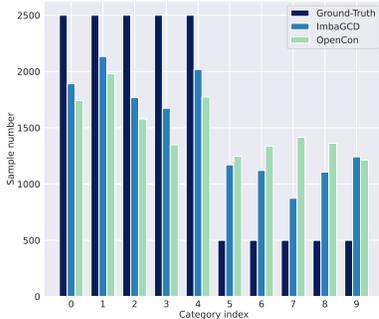


Figure 3. Comparison of the predicted sample numbers for each class on the CIFAR-10 dataset, using a step decreasing strategy with an imbalance factor (ρ) of 5.

a batch size of 512. We employ stochastic gradient descent with a momentum of 0.9 and a weight decay of $10e-4$. The learning rate commences at 0.02 and undergoes decay by a factor of 10 at the 50% and 75% stages of the training process. The momentum for updating the prototype and class prior, represented as μ , is consistently maintained at 0.99.

4.2. Compared with SOTA

Baselines In our study, we evaluate our approach against three state-of-the-art GCD methods serving as baselines: (1) GCD [44], the pioneering framework for GCD, efficiently identifies and clusters unknown categories in high-dimensional data without exhaustive labeling; (2) ORCA [4], combines supervised and unsupervised learning to effectively utilize labeled and unlabeled data, addressing class imbalance in open-world settings; (3) OpenCon [41], employs contrastive learning to maximize similarity between positive pairs while minimizing it for negative pairs, enhancing adaptability to new categories with minimal super-

Table 3. Ablation study on loss component in CIFAR100. Known class, unknown-aware(Un1), and unknown-agnostic (Un2) accuracies.

IMF	Loss	Known	Un1	Un2
$\rho = 0.5$	w/o L_{sup}	42.86	46.51	46.91
	w/o L_{ins}	64.96	17.40	12.09
	w/o L_{proto}	68.36	46.15	44.21
	Ours	68.09 ± 0.13	47.92 ± 0.33	46.22 ± 0.33
$\rho = 1$	w/o L_{sup}	43.20	42.90	39.96
	w/o L_{ins}	63.68	17.54	6.43
	w/o L_{proto}	67.57	42.91	38.27
	Ours	67.98 ± 0.37	43.39 ± 0.59	38.76 ± 0.90
$\rho = 5$	w/o L_{sup}	31.01	30.40	9.25
	w/o L_{ins}	52.23	18.24	12.07
	w/o L_{proto}	66.80	32.72	21.20
	Ours	67.82 ± 0.06	37.87 ± 1.59	27.64 ± 2.26
$\rho = 10$	w/o L_{sup}	30.33	29.92	4.50
	w/o L_{ins}	51.75	20.88	13.42
	w/o L_{proto}	66.52	32.56	14.76
	Ours	67.82 ± 0.07	34.87 ± 0.70	21.68 ± 0.29

vision. We evaluate our method against the aforementioned baselines in various scenarios: original GCD setting, balanced setting, and multiple imbalanced settings, allowing for a comprehensive comparison across diverse conditions.

ImbaGCD achieves SOTA performance In Tables 1 and 2, ImbaGCD demonstrates a significant performance advantage over its competitors on both CIFAR-100 and ImageNet datasets, particularly in the novel classes. Specifically, on the CIFAR-100 dataset, our method achieves an improvement of approximately 2 - 4% over the best baseline in terms of both unknown-aware and unknown-agnostic accuracy across a range of imbalanced settings. Regarding the ImageNet-100 dataset, ImbaGCD surpasses the baseline performance by approximately 14-16% and 15-19% under $\rho = 5$ and $\rho = 10$ settings, respectively. Additionally, it attains competitive results in balanced and original GCD settings ($\rho = 0.5$). It is important to highlight that the unknown-agnostic evaluation presents a greater challenge compared to the unknown-aware evaluation, particularly in the context of highly imbalanced settings. Our method exhibits more substantial improvements under these challenging evaluation conditions, with the enhancements in unknown-agnostic performance being more prominent than those in unknown-aware performance.

4.3. Ablation Study

Class distribution prediction ImbaGCD outperforms the state-of-the-art OpenCon in predicting class distri-

butions for different decreasing strategies on the CIFAR10. We conduct experiments on two decreasing types: exponential-decreasing (Figure 2) and step-decreasing (Figure 3). The x -axis represents the class index, while the y -axis denotes the sample number of each class. Upon comparing our results with those of OpenCon [41] and the ground-truth sample numbers, it becomes evident that our method consistently attains a superior class distribution. Furthermore, the deviations between the predictions made by our method and the ground-truth values are consistently smaller than those observed for OpenCon across all classes. These findings underscore the effectiveness and robustness of our approach in predicting class distributions.

Analysis of the loss components Recall our objective function in Eq. 11 has three components. We perform an ablation study (Table 3, with Un1 as Un1 accuracy and Un2 as Un2 accuracy) to analyze their contributions. The ImbaGCD model is modified by removing: L_{sup} , L_{ins} , and L_{proto} . This study aims to understand each component’s impact on performance. We make several observations from our ablation study:

- When there are more unknown class samples (e.g., $\rho = 0.5, 1$), removing L_{sup} mainly affects known class accuracy with less impact on Un1 and Un2 ($< 1\%$). For fewer unknown samples (e.g., $\rho = 5, 10$), reductions in known class accuracy cause significant drops in Un1 ($\sim 5 - 7\%$) and Un2 ($17 - 18\%$).
- For L_{ins} , at $\rho = 0.5$ and $\rho = 1$, the known class accuracy decreases slightly, while the declines in Un1 and Un2 are more pronounced ($\sim 26 - 32\%$). However, when $\rho = 5$ and $\rho = 10$, there is a noticeable drop in all metrics (known, Un1, and Un2)
- For L_{proto} , the effects are not significant when $\rho = 0.5$ and $\rho = 1$. However, there is a notable improvement in Un1 ($\sim 2 - 5\%$) and Un2 ($\sim 6 - 7\%$) when the L_{proto} component is included.

5. Related Work

Generalized category discovery (GCD) This problem extends NCD [14, 17, 18, 32] by considering unlabeled data from both known and novel classes [44]. GCD addresses this challenge through semi-supervised contrastive learning on large-scale pre-trained visual transformers (ViT) followed by constraint KMeans [1]. Concurrently, ORCA [4] proposes an uncertainty adaptive margin loss to reduce intra-class variances between known and novel classes. Opencon [41] proposes a contrastive learning frameworks which selects the positive and negative pair via a moving average prototype. Despite the prevalence of class imbalance, many existing works in this area overlook its impact. Our work contributes to the field by addressing this gap.

Learning with class-imbalanced data Real-world datasets often exhibit a long-tailed label distribution [34, 43], complicating standard DNN training and generalization [13, 36, 47]. To address class imbalance, approaches include (a) re-weighting loss functions class-wise [5, 30, 35], and (b) re-sampling datasets for balanced training distribution [3, 8]. Both methods perform better when applied in later training stages for DNNs [23, 42]. However, they assume full supervision, prompting studies on weak-supervision, such as semi-supervised learning [26, 48], where [26] requires ground-truth class priors and [48] estimates them from labeled data. Our work, however, presents a greater challenge as the imbalances occur between known and unknown classes, with no prior class distribution information available to estimate the known class prior distribution.

6. Conclusion

In this paper, we present the significant and challenging problem of Imbalanced Generalized Category Discovery (ImbaGCD), characterized by an imbalanced distribution of unlabeled data. To tackle this issue, we develop ImbaGCD, a novel and robust optimal transport-based expectation maximization framework. Our extensive experimental evaluation encompasses varying settings, including balanced and multiple imbalanced scenarios. The results demonstrate that our proposed method consistently outperforms state-of-the-art approaches across diverse imbalanced settings. Furthermore, ImbaGCD exhibits competitive performance in both balanced and original GCD settings, highlighting its adaptability and effectiveness across a range of situations. These findings establish ImbaGCD as a highly capable and versatile solution for addressing the ImbaGCD problem, paving the way for further advancements in GCD.

References

- [1] David Arthur and Sergei Vassilvitskii. K-means++ the advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035, 2007. 8
- [2] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018. 2
- [3] Jonathon Byrd and Zachary Lipton. What is the effect of importance weighting in deep learning? In *International conference on machine learning*, pages 872–881. PMLR, 2019. 8
- [4] Kaidi Cao, Maria Brbic, and Jure Leskovec. Open-world semi-supervised learning. In *ICLR*, 2020. 2, 4, 5, 7, 8
- [5] Kaihua Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in Neural Information Processing Systems*, 32:1565–1576, 2019. 2, 8

- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020. 4
- [7] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009. 1
- [8] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002. 2, 8
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pages 1597–1607. PMLR, 2020. 4, 5
- [10] Haoang Chi, Feng Liu, Wenjing Yang, Long Lan, Tongliang Liu, Bo Han, Gang Niu, Mingyuan Zhou, and Masashi Sugiyama. Meta discovery: Learning to discover novel classes given very limited data. In *ICLR*, 2021. 2
- [11] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9268–9277, 2019. 2
- [12] Marco Cuturi. Sinkhorn distances: lightspeed computation of optimal transport. In *Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 2*, pages 2292–2300, 2013. 2, 3, 4
- [13] Qi Dong, Shaogang Gong, and Xiatian Zhu. Imbalanced deep learning by minority class incremental rectification. *IEEE transactions on pattern analysis and machine intelligence*, 41(6):1367–1381, 2018. 8
- [14] Enrico Fini, Enver Sangineto, Stéphane Lathuilière, Zhun Zhong, Moin Nabi, and Elisa Ricci. A unified objective for novel class discovery. In *ICCV*, 2021. 1, 8
- [15] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017. 5
- [16] Mohamed Farouk Abdel Hady and Friedhelm Schwenker. Semi-supervised learning. *Handbook on Neural Information Processing*, pages 215–239, 2013. 1
- [17] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Autonovel: Automatically discovering and learning novel visual categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 1, 8
- [18] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *CVPR*, 2019. 1, 8
- [19] Haibo He and Eduardo A Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009. 2
- [20] 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. 4
- [21] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5375–5384, 2016. 2
- [22] Bingyi Kang, Saining Li, and Dacheng Tao. Decoupling representation and classifier for long-tailed recognition. In *Eighth International Conference on Learning Representations (ICLR)*, 2020. 2
- [23] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019. 8
- [24] Salman H Khan, Munawar Hayat, Fatih Porikli, and Mohammed Bennamoun. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 29(8):3573–3587, 2018. 2
- [25] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 4
- [26] Jaehyung Kim, Youngbum Hur, Sejun Park, Eunho Yang, Sung Ju Hwang, and Jinwoo Shin. Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. *Advances in neural information processing systems*, 33:14567–14579, 2020. 8
- [27] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [28] Junnan Li, Caiming Xiong, and Steven CH Hoi. Mopro: Webly supervised learning with momentum prototypes. *arXiv preprint arXiv:2009.07995*, 2020. 5
- [29] Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*, 2020. 4
- [30] Mingchen Li, Xuechen Zhang, Christos Thrampoulidis, Jiasi Chen, and Samet Oymak. Autobalance: Optimized loss functions for imbalanced data. *Advances in Neural Information Processing Systems*, 34:3163–3177, 2021. 8
- [31] Ziyun Li, Jona Otholt, Ben Dai, Christoph Meinel, Haojin Yang, et al. Supervised knowledge may hurt novel class discovery performance. *Transactions of Machine Learning Research*. 2
- [32] Ziyun Li, Jona Otholt, Ben Dai, Christoph Meinel, Haojin Yang, et al. A closer look at novel class discovery from the labeled set. *arXiv preprint arXiv:2209.09120*, 2022. 1, 2, 8
- [33] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019. 2
- [34] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Barambe,

- and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pages 181–196, 2018. 8
- [35] Seulki Park, Jongin Lim, Younghan Jeon, and Jin Young Choi. Influence-balanced loss for imbalanced visual classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 735–744, 2021. 8
- [36] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pages 4334–4343. PMLR, 2018. 8
- [37] Mamshad Nayeem Rizve, Navid Kardan, Salman Khan, Fahad Shahbaz Khan, and Mubarak Shah. Openldn: Learning to discover novel classes for open-world semi-supervised learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pages 382–401. Springer, 2022. 2
- [38] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boult. Toward open set recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(7):1757–1772, 2012. 1
- [39] Richard Sinkhorn. Diagonal equivalence to matrices with prescribed row and column sums. ii. *Proceedings of the American Mathematical Society*, 1974. 3
- [40] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. 1
- [41] Yiyou Sun and Yixuan Li. Opencon: Open-world contrastive learning. *Transactions of Machine Learning Research*. 2, 4, 5, 7, 8
- [42] Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *Advances in Neural Information Processing Systems*, 33:1513–1524, 2020. 8
- [43] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018. 8
- [44] 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*, 2022. 1, 2, 5, 7, 8
- [45] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008. 2
- [46] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009. 3
- [47] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. *Advances in neural information processing systems*, 30, 2017. 8
- [48] Chen Wei, Kihyuk Sohn, Clayton Mellina, Alan Yuille, and Fan Yang. Crest: A class-rebalancing self-training framework for imbalanced semi-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10857–10866, 2021. 8