

Dissecting Generalized Category Discovery: Multiplex Consensus under Self-Deconstruction

Luyao Tang^{1,2,5*}, Kunze Huang^{1,2,*}, Chaoqi Chen^{3,†}, Yuxuan Yuan¹, Chenxin Li⁴

Xiaotong Tu^{1,2}, Xinghao Ding^{1,2} and Yue Huang^{1,2}

¹ Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University ² School of Informatics, Xiamen University

³ Shenzhen University ⁴ The Chinese University of Hong Kong ⁵ The University of Hong Kong

{lytang, kzhuang}@stu.xmu.edu.cn, cqchen1994@gmail.com, yhuang2010@xmu.edu.cn

Abstract

Human perceptual systems excel at inducing and recognizing objects across both known and novel categories, a capability far beyond current machine learning frameworks. While generalized category discovery (GCD) aims to bridge this gap, existing methods predominantly focus on optimizing objective functions. We present an orthogonal solution, inspired by the human cognitive process for novel object understanding: decomposing objects into visual primitives and establishing cross-knowledge comparisons. We propose **ConGCD**, which establishes primitive-oriented representations through high-level semantic reconstruction, binding intra-class shared attributes via deconstruction. Mirroring human preference diversity in visual processing, where distinct individuals leverage dominant or contextual cues, we implement dominant and contextual consensus units to capture class-discriminative patterns and inherent distributional invariants, respectively. A consensus scheduler dynamically optimizes activation pathways, with final predictions emerging through multiplex consensus integration. Extensive evaluations across coarse- and fine-grained benchmarks demonstrate ConGCD’s effectiveness as a consensus-aware paradigm. *Code is here.*

1. Introduction

Generalizing inductive and recognition capabilities to unknown categories of objects is an innate ability of human vision [8], but it is incredibly challenging for machine learning models. Generalized Category Discovery (GCD) [39, 48, 52, 53, 57] is an advanced method to address this issue, which encourages models to cluster and

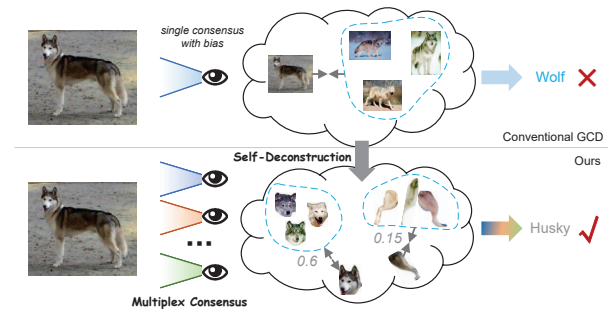


Figure 1. Generalized category discovery inherently demands multiplex consensus via self-deconstruction, rejecting image-level processing in favor of attribute-aware compositional reasoning.

recognize a subset of samples within the data that includes both unlabeled and unknown categories as well as known categories after training on partially labeled data.

From the perspective of representation learning [44], current GCD paradigms including contrastive learning [10, 48], prototype classifiers [53], and sample selection [31, 39] share a fundamental oversight: they process images as atomic entities, neglecting the compositional nature of visual recognition. For instance, CMS [10] applies mean-shift to global embeddings, while SimGCD [53] distills category prototypes from holistic representations. Though effective for coarse discrimination, these methods fail to emulate the human capacity for attribute-level reasoning - the ability to decompose objects into semantic primitives and establish cross-category consensus through comparative analysis.

Humans can transfer previously acquired knowledge when learning new concepts [22, 55]. If young children learn to distinguish between “cows” and “sheep” based on their shape, they can apply the same rule to differentiate between “pigs” and “horses”. However, when it comes to telling “pandas” and “bears” apart, relying solely on shape can lead to confusion like Fig. 1; attention to additional

* Equal Contribution; † Corresponding Author.

visual primitives [16] such as stripes and ears is required. For instance, the Oxford Dictionary [13] defines a giraffe as “a tall animal with a very long neck, long legs, and dark marks on its coat”. This definition involves two shape primitives (long neck, long legs) and one texture-color attribute (dark marks). Unfortunately, shortcut learning [19] in deep networks typically activates a few primitives (such as the gray skin of an elephant [18]), causing confusion between categories with similar primitives, like rhinoceroses.

In this work, we propose cognitive science-inspired ConGCD, which attempts to address the above challenge by focusing on two aspects as shown in Fig. 2: the **self-deconstruction of visual primitives** and the **formation of multiplex consensus**. (1) For the decoupling of visual primitives, we define a set of primitive-oriented representations to bind intra-class shared visual primitives to spatial regions [8] in a construction manner. The competitive relationship among primitive-oriented representations ensures that each primitive mask focuses on a single visual attribute. Collectively, all primitive-oriented representations can be seen as a synthesis of the image’s key visual primitives. (2) For the formation of multiplex consensus, unlike monolithic feature aggregation, we implement dominant and contextual consensus units: dominant units retain high-response neurons to capture class-discriminative patterns, while contextual units preserve weaker activations encoding distributional invariants. A dynamic consensus scheduler optimizes activation pathways between these units, with final predictions emerging through multiplex integration of their complementary perspectives. This process mimics how humans iteratively reconcile dominant cues (e.g., beak shape for birds) with contextual invariants (e.g., flight posture) when categorizing novel exemplars [24, 25].

Our main contributions can be summarized as follows:

- A mechanism inducing visual primitives via competitive binding, enabling unsupervised discovery of primitive-discriminative attributes through topological competition.
- A multiplex consensus framework: dominant units retain high-response neurons for class-related variables, while contextual units preserve weak activations for inherent distributions, coordinated by dynamic scheduling.
- Extensive validation showing our paradigm’s plug-and-play compatibility with GCD methods, achieving state-of-the-art performance across seven benchmarks.

2. Related Works

2.1. Generalized Category Discovery

In Generalized Category Discovery (GCD) [48], the goal is to classify unlabeled samples into known and unknown classes using a limited labeled dataset from known classes, reflecting a realistic approach to real-world image recognition. Initially, GCD was formulated by [48], employing

contrastive and semi-supervised learning methods that relied heavily on clustering during inference. Consequently, SimGCD, a robust baseline approach that replaces clustering with a classifier, is trained using a pseudo-labeling strategy that has shown remarkable effectiveness in SSL [53]. To address this limitation, subsequent works aim to improve feature representation by exploiting underlying cross-instance relationships [37]. An EM-like framework is proposed to alternate between contrastive learning and class number estimation [57]. Recent works have proposed using parametric classifiers to avoid prediction biases [14]. However, previous studies [31, 34] focused on class tokens for clustering, overlooking that similar visual attributes across samples can cause inter-class boundary confusion.

2.2. Object-Centric Learning

In the evolving landscape of Object-Centric Learning (OCL), Slot Attention established a significant milestone [30]. This mechanism enabled further research, like [14], which extended OCL to videos, using temporal consistency and dynamics within video sequences to better differentiate objects from backgrounds. Furthermore, to address the stability issue in single-view OCL, [21] introduces modules to guide slots away from background noise and towards consistent object representation. [2] explores self-supervision, offering the first unsupervised multi-object segmentation in real-world videos. And the challenge of adapting OCL to complex real-world settings was tackled by [41], bridging the gap between lab and real-world scenarios. OCL currently concentrates on object-level separation [28], and we extend it to a more fine-grained visual attribute discovery level, and let the model innately have the inductive ability for the semantics of visual attributes.

3. Preliminaries

3.1. Visual Primitives

The generation process of a sample \mathbf{x} is actually a combination process of a set of visual primitives s^1, \dots, s^K ($s^{1:K}$) with $s^i \in \mathcal{S}$, where \mathcal{S} is a finite set [16]. Considering the joint distribution of \mathbf{x} and $s^{1:K}$, due to biases in observation and annotation and deviation from the real-world distribution, the label of a sample can depend on one or more primitives. The data generation process can be expressed as

$$\begin{aligned} \mathbf{z} &\sim p(\mathbf{z}), \quad s^i \sim p(s^i|\mathbf{z}), \quad \mathbf{x} \sim p(\mathbf{x}|\mathbf{z}), \\ p(s^{1:K}, \mathbf{x}) &= p(s^{1:K}) \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z}|s^{1:K})d\mathbf{z}, \end{aligned} \quad (1)$$

where \mathbf{z} is a factor in the latent space. We believe that if different marginal distributions of given primitives exist but they share the same conditional generation process, there will be distribution shift between categories, that is

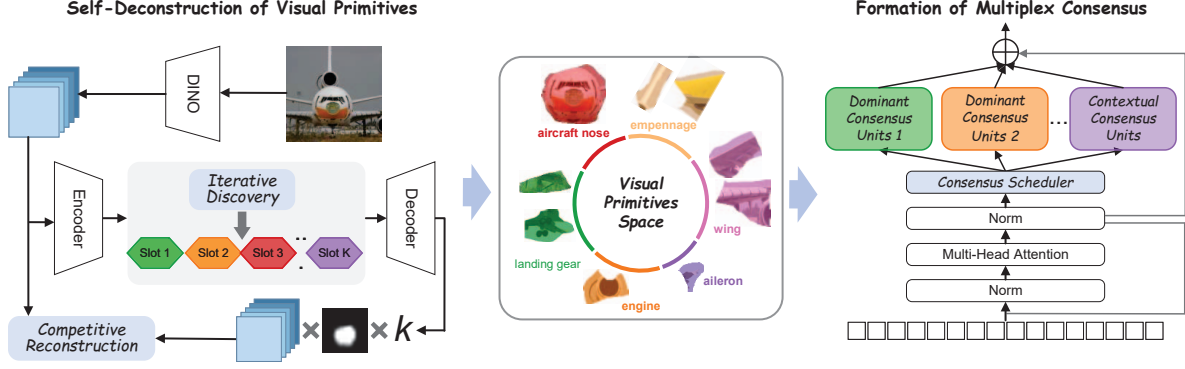


Figure 2. Overview of the proposed ConGCD: (1) Visual primitives are discovered iteratively via competitive reconstruction; (2) Each visual primitive is processed by dominant consensus units with primitive mask, while the inherent distribution is processed by contextual consensus units scheduler with holistic representation. The weight of each unit is determined by the consensus scheduler.

$p_{\text{class-}n}(s^{1:K}) \neq p_{\text{class-}n'}(s^{1:K})$. However, due to the shared generative model between distributions, certain primitives subset $\tilde{\mathcal{A}}$ are highly similar across inter-class, we have $p_{\text{class-}n}(s^{1:K}, \mathbf{x}) = p_{\text{class-}n'}(s^{1:K}) \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z}|\tilde{s}^{1:K})d\mathbf{z}$, where $\tilde{s}^i \in \tilde{\mathcal{S}}$ and $\tilde{\mathcal{S}} \subseteq \mathcal{S}$. The similarity of inter-class primitives is disastrous for distinguishing between categories. Therefore, we aim to find a complementary set $\mathcal{S} \setminus \tilde{\mathcal{S}}$, the discriminative visual primitives that maximize the distribution difference between categories, serving as the factors that contribute the most to inter-class separation.

3.2. Problem Definition of GCD

we consider a labeled subset for each dataset, denoted as $\mathcal{D}_l = \{(\mathbf{x}_i^l, y_i^l)\} \subset \mathcal{X} \times \mathcal{Y}_l$. Additionally, there is an unlabeled subset $\mathcal{D}_u = \{(\mathbf{x}_i^u, y_i^u)\} \subset \mathcal{X} \times \mathcal{Y}_u$. In \mathcal{D}_l , only known classes are present, i.e., $\mathcal{Y}_l = \mathcal{C}_{\text{known}}$. On the other hand, \mathcal{D}_u contains both known and novel classes, $\mathcal{Y}_u = \mathcal{C}_{\text{known}} \cup \mathcal{C}_{\text{novel}}$. The task for models is to perform clustering on both the known and novel classes in \mathcal{D}_u . The number of novel classes, represented as N_{novel} , can be determined in advance [36, 48, 57]. The functions $f(\cdot)$ and $g(\cdot)$ act as the feature extractor and projection head respectively. Both the feature $\mathbf{h}_i = f(\mathbf{x}_i)$ and the projected embedding $\mathbf{z}_i = g(\mathbf{h}_i)$ are subjected to L-2 normalization.

4. Methodology

4.1. Self-Deconstruction of Visual Primitives

Visual primitives Iterative Discovery. Referring to the process in which humans compare the visual features of an object with the existing knowledge in the brain and conduct cyclic searches after receiving visual signals [8], we define visual primitive discovery as an iterative process. It originates from slot-attention [30]. We predefine primitive-centric representations (learnable embeddings), $\mathbf{s} \sim \mathcal{N}(\mathbf{s}; \boldsymbol{\mu}, \boldsymbol{\sigma}) \in \mathbb{R}^{K \times D_s}$. K is the number of visual

primitives, and D_s is the number of dimensions.

For the pre-trained encoder's output $\mathbf{z} = g(f(\mathbf{x})) \in \mathbb{R}^{N \times D_z}$, the attention mechanism of the human-like perception process [46] is implemented. Representing \mathcal{K}_β , \mathcal{Q}_γ and \mathcal{V}_ϕ for \mathbf{z} as the projection networks of key, query, and value with parameters β , γ , and ϕ , respectively. The discovery process of visual primitives $\text{vpdis}(\cdot)$ and the attention function $\text{attn}(\cdot)$ can be defined as

$$\begin{aligned} \text{vpdis}(\mathbf{A}, \mathbf{v}) &= \mathbf{A}^T \mathbf{v}, \quad A_{ij} = \frac{\text{attn}(\mathbf{q}, \mathbf{k})_{ij}}{\sum_{l=1}^K \text{attn}(\mathbf{q}, \mathbf{k})_{lj}}, \\ \text{attn}(\mathbf{q}, \mathbf{k}) &= \frac{e^{M_{ij}}}{\sum_{l=1}^N e^{M_{il}}}, \quad \mathbf{M} = \frac{\mathbf{k}\mathbf{q}^T}{\sqrt{D_s}}, \end{aligned} \quad (2)$$

Where \mathbf{A} is the attention matrix. The $\mathbf{k} = \mathcal{K}_\beta(\mathbf{z}) \in \mathbb{R}^{N \times D_s}$ and $\mathbf{v} = \mathcal{V}_\phi(\mathbf{z}) \in \mathbb{R}^{N \times D_s}$ are the key and value vectors, respectively. To ensure the binding of the attention relationship between visual primitives and feature embeddings, queries are a function of visual primitives \mathbf{s} , and the formal expression is $\mathbf{q} = \mathcal{Q}_\gamma(\mathbf{s}) \in \mathbb{R}^{K \times D_s}$.

Similar to the process of thinking and reasoning [1], visual primitives are iteratively refined over T attention iterations. Specifically, in iteration t , the queries are given by $\hat{\mathbf{q}}^t = \mathcal{Q}_\gamma(\mathbf{s}^t)$. The Gated Recurrent Unit (GRU) [11, 12] is applied, denoted as \mathcal{G}_θ . The visual primitives update process can be summarized as $\mathbf{s}^{t+1} = \text{vpdis}(\text{attn}(\hat{\mathbf{q}}^t, \mathbf{k}), \mathbf{v})$.

Given \mathbf{x} at $t = T$ iteration, the posterior distribution of the visual primitives is

$$p(\mathbf{s}^T | \mathbf{x}) = \delta \left(\mathbf{s}^T - \prod_{t=1}^T \mathcal{G}_\theta(\mathbf{s}^{t-1}, \text{vpdis}(\text{attn}(\hat{\mathbf{q}}^{t-1}, \mathbf{k}), \mathbf{v})) \right), \quad (3)$$

where $\delta(\cdot)$ represents the Dirac delta distribution [4], we sample \mathbf{s} from a Gaussian distribution [20], at $t = 0$ itera-

tion for initialization so that the distribution of \mathbf{s}^0 represents relatively uniform visual primitives, preventing biases in the visual primitive discovery process.

Deconstruction via Competitive Reconstruction. $\hat{\mathbf{z}} = \Phi_d(\mathbf{s}_k)$ The discovery process of visual primitives is unsupervised, and we complete it through the surrogate task of reconstruction. After \mathbf{s}^T goes through the above iterative discovery, it is sent in as a general lightweight multi-layer perceptron (MLP) as the decoder [42], denoted as $\Phi_d(\cdot)$. $\Phi_d(\cdot)$ spatially broadcasts K visual primitives to the same spatial dimension as $\hat{\mathbf{z}}_k = \Phi_d(\mathbf{s}_k) \in \mathbb{R}^{N \times (D_z + 1)}$. The reconstructed $\hat{\mathbf{z}}_k$ is composed of $D_z + 1$ channels. After channel division, we obtain the reconstructed feature embeddings $\hat{\mathbf{z}}_k \in \mathbb{R}^{N \times D_s}$ and the activation region $\alpha_k \in \mathbb{R}^{N \times 1}$. In order to bind the visual primitive embedding to the region with the same semantics in the image, we use softmax on α_k for normalization over K visual primitives to introduce the competitive relationship between visual primitives, so that a certain primitive is associated with certain pixels

$$\hat{\mathbf{z}} = \sum_{k=1}^K \hat{\mathbf{z}}_k \odot \mathbf{m}_k, \quad \mathbf{m}_k = \text{softmax}_k(\alpha_k). \quad (4)$$

The optimization objective of the iterative discovery and decoupling process of visual primitives is simple

$$\mathcal{L}_{\text{rec}} = \|\hat{\mathbf{z}} - \mathbf{z}\|^2, \quad \hat{\mathbf{z}} = \Phi_d(\mathbf{s}). \quad (5)$$

This self-supervised deconstruction establishes basic visual primitives through competitive binding, enabling machines to mirror human attribute-level perception. By disentangling objects into composable factors, we equip models with intrinsic interpretability for fine-grained comparisons while maintaining adaptability to novel category structures.

4.2. Formation of Multiplex Consensus

Models need to understand the visual primitives. However, exclusive reliance on dominant feature activations [18] creates self-reinforcing biases that degrade unknown-class generalization, as shown in Tab. 4, where strongly activated neurons encode class-agnostic structural primitives, while their weakly activated counterparts preserve cross-category contextual patterns [29].

This duality (as visualized in Fig. 3) drives our consensus formation through two complementary pathways: (1) Dominant Consensus Units distill class-discriminative knowledge by routing high-activation visual primitives through neural competition layers, preventing primitive overcommitment to known categories. (2) Contextual Consensus Units continuously assimilate weakly activated embeddings to model class-transcendent relationships. The resultant multiplex architecture achieves self-stabilization through consensus scheduler, where primitive specificity and contextual generality co-evolve.




Figure 3. Dominant and weak neurons capture the class-related variables and inherent distributions, respectively.

Dominant Consensus Units. For the output $\mathbf{z}^i \in \mathbb{R}^{N \times D_z^i}$ of the i -th block in the network forward process, in order to be consistent with the pixel-level relationship of visual primitives, we selectively retain the top r (%) of neuron activation responses at the token-level from the spatial dimension. Specifically, we customize a sorting function $\text{sort}(\cdot)$ from large to small for neurons to generate \mathbf{z}_{attr}^i after filtering neuron activation values. Furthermore, for each consensus units, their attention should be focused on a single visual primitive to ensure the expertise of the units (Analogous to human visual preferences, some people rely on the head to distinguish animals, and some people rely on the tail). We directly use \mathbf{m}_k to represent the activation of the visual primitive region to limit the input of k -th unit. The forward process of each dominant consensus unit is

$$\mathbf{z}_{Dom}^{ik} := \begin{cases} \mathbf{z}^i \odot \mathbf{m}_k, & \text{if } \mathbf{z}^i[:, :] \geq \text{Top-R}_k(\text{sort}(\mathbf{z}^i[:, :])) \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $R_k = N \times r_k$ in $\text{Top-R}_k(\cdot)$, for visual primitive schedulers used to capture class-related variables, they need to perceive highly activated neurons and provide features conducive to separation between classes.

Contextual Consensus Units. Due to the inherent distribution of the dataset, only maximizing the extraction of class-related visual primitives will cause difficulties for the model in recognizing novel classes. Therefore, we propose contextual consensus units based on weak neurons. It accepts the complete \mathbf{z}^i as input and perceives all visual primitives of the object. We decompose the low-activation region in \mathbf{z}^i

$$\mathbf{z}_{Con}^i := \begin{cases} \mathbf{z}^i, & \text{if } \mathbf{z}^i[:, :] \leq \text{Top-R}(\text{sort}(\mathbf{z}^i[:, :])) \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

which is similarly to Eq. (6), but dominant consensus units on class-independent features activated in weak neurons, aiming to extract the information common to each sample from the feature embedding itself. This prevents the model

	Method	CUB-200			FGVC-Aircraft			Stanford-Cars			Average		
		All	Known	Novel	All	Known	Novel	All	Known	Novel	All	Known	Novel
DINOv1	GCD [48]	51.3	56.6	48.7	45.0	41.1	46.9	39.0	57.6	29.9	45.1	51.8	41.8
	GPC [57]	52.0	55.5	47.5	43.3	40.7	44.8	38.2	58.9	27.4	44.5	51.7	39.9
	XCon [15]	52.1	54.3	51.0	47.7	44.4	49.4	40.5	58.8	31.7	46.8	52.5	44.0
	PromptCAL [56]	62.9	64.4	62.1	52.2	52.2	52.3	50.2	70.1	40.6	55.1	62.2	51.7
	AMEND [3]	64.9	75.6	59.6	52.8	61.8	48.3	56.4	73.3	48.2	58.0	70.2	52.0
	μ GCD [49]	65.7	68.0	64.6	53.8	55.4	53.0	56.5	68.1	50.9	58.7	63.8	56.2
	CMS [10]	68.2	76.5	64.0	56.0	63.4	52.3	56.9	76.1	47.6	60.4	72.0	54.6
	InfoSieve [38]	69.4	77.9	65.2	56.3	63.7	52.5	55.7	74.8	46.4	60.5	72.1	54.7
	SimGCD [53]	60.3	65.6	57.7	54.2	59.1	51.8	53.8	71.9	45.0	56.1	65.5	51.5
	+ ConGCD	61.6	65.1	59.5	55.0	58.4	53.1	54.5	72.2	47.8	57.1	65.2	53.5
	LegoGCD [6]	63.8	71.9	59.8	55.0	61.5	51.7	57.3	75.7	48.4	58.7	69.7	53.3
	+ ConGCD	65.3	72.6	61.7	55.6	61.2	54.0	58.1	75.9	49.5	59.7	69.9	55.1
	SPTNet [52]	65.8	68.8	65.1	59.3	61.8	58.1	59.0	79.2	49.3	61.4	69.9	57.5
	+ ConGCD	68.1	68.5	67.8	59.7	61.3	59.2	59.1	<u>79.0</u>	<u>49.8</u>	62.3	69.6	58.9
	SelEx [39]	78.7	81.3	<u>77.5</u>	60.9	70.3	56.2	57.0	77.4	47.2	<u>65.5</u>	76.3	<u>60.3</u>
+ ConGCD	81.7	<u>80.4</u>	82.4	62.5	<u>70.2</u>	<u>58.7</u>	57.5	77.5	47.9	67.3	<u>76.0</u>	63.0	
Avg. Δ	+2.03	-0.25	+2.83	+0.85	-0.40	+1.80	+0.53	+0.10	+1.28	+1.18	-0.18	+1.98	
DINOv2	GCD [48]	71.9	71.2	72.3	55.4	47.9	59.2	65.7	67.8	64.7	64.3	62.3	65.4
	SimGCD[53]	71.5	78.1	68.3	63.9	69.9	60.9	71.5	81.9	66.6	69.0	76.6	65.3
	μ GCD [49]	74.0	75.9	73.1	66.3	68.7	65.1	76.1	91.0	68.9	72.1	78.5	69.0
	SelEx[39]	86.0	<u>86.5</u>	<u>85.7</u>	82.1	84.8	<u>80.7</u>	80.5	<u>91.8</u>	75.0	82.9	<u>87.7</u>	80.5
	+ ConGCD	86.3	87.4	85.8	<u>81.7</u>	<u>83.3</u>	81.0	<u>79.8</u>	93.1	<u>73.3</u>	<u>82.6</u>	87.9	<u>80.1</u>

Table 1. Results on the **fine-grained** semantic shift benchmark. ConGCD offers excellent cross-scheme and cross-model compatibility.

from overly focusing on the visual primitives of known classes and from effectively perceiving novel classes.

Consensus Scheduler. The self-deconstruction process manifests in ViT blocks [7] through dual activation polarization: while multi-head self-attention (MHA) [46] captures cross-patch dependencies, the decomposed dominant/contextual consensus pathways exhibit complementary neuron activation patterns. We reformulate the FFN as parallel consensus units governed by activation separation:

$$\begin{aligned}
 f_{\text{ConGCD}}(\mathbf{z}) &= \sum_{k=1}^{K+M} G(\mathbf{z})_k \cdot E_k(\mathbf{z}_{\text{Dom}}^k / \mathbf{z}_{\text{Con}}) \\
 &= \sum_{k=1}^{K+M} \text{softmax}(\mathbf{W}\mathbf{z}) \cdot \mathbf{W}_{\text{FFN}_k}^2 \phi(\mathbf{W}_{\text{FFN}_k}^1 \mathbf{z}_{\text{Dom}}^k / \mathbf{z}_{\text{Con}}),
 \end{aligned} \tag{8}$$

where \mathbf{W} and $\mathbf{W}_{\text{FFN}_k}^{1/2}$ are the learnable parameters. $\phi(\cdot)$ is a nonlinear activation function. The primitive scheduler operates densely for K dominant units and M contextual units and assigns a non-zero probability to each unit.

Due to inconsistent preferences across units, determining the ratio r_k (%) for detecting dominant/contextual activation neurons becomes challenging. We address this via an Activation Separation Distributor (ASR) that dynamically generates unit-specific separation ratios from \mathbf{z} , enabling adaptive boundary determination for activation decompo-

sition per sample. Scheduler adopt a new R^* to screen neurons using $\text{Top-R}_k^*(\cdot)$:

$$R_k^* = N \times (r_k + \eta \cdot \text{softmax}(\mathbf{W}_{\text{ASR}}\mathbf{z})[k]), \tag{9}$$

where \mathbf{W}_{ASR} is learnable, η controls scaling, and N denotes patch count. This allows sample-aware adaptation to diverse visual primitives while maintaining activation balance between classes. In general, given \mathbf{z} as the input, the update process is

$$\mathbf{z}_{\text{out}} = f_{\text{ConGCD}}(\text{LN}(\mathbf{z})) + \mathbf{z}, \tag{10}$$

where $\text{LN}(\cdot)$ indicates layer normalization, and \mathbf{z}_{out} is the output and is sent to the next block. The multiplex consensus framework resolves fundamental tension in GCD through symbiotic coordination, dominant units sharpen categorical boundaries while contextual units maintain open-world continuity, mirroring human vision’s dual processing of focal attributes and environmental context.

5. Experiments

Through comprehensive experimental verification and analysis, we aim to answer the following questions. (1) Can ConGCD improve the accuracy of GCD? (2) Can ConGCD be plugged and played in existing GCD schemes? (3) What is the source of the performance improvement of ConGCD? (4) Is ConGCD robust to changes in hyperparameters?

	Method	CIFAR-10			CIFAR-100			ImageNet-100			Average		
		All	Known	Novel	All	Known	Novel	All	Known	Novel	All	Known	Novel
DINOv1	ORCA [5]	96.9	95.1	97.8	74.2	82.1	67.2	79.2	93.2	72.1	83.4	90.1	79.0
	GCD [48]	91.5	97.9	88.2	73.0	76.2	66.5	74.1	89.8	66.3	79.5	88.0	73.7
	GPC [57]	90.6	97.6	87.0	75.4	84.6	60.1	75.3	93.4	66.7	80.4	91.9	71.3
	XCon [15]	96.0	97.3	95.4	74.2	81.2	60.3	77.6	93.5	69.7	82.6	90.7	75.1
	PIM [9]	94.7	97.4	93.3	78.3	84.2	66.5	83.1	95.3	77.0	85.4	92.3	78.9
	PromptCAL [56]	97.9	96.6	<u>98.5</u>	81.2	84.2	75.3	83.1	92.7	78.3	87.4	91.2	84.0
	DCCL [36]	96.3	96.5	96.9	75.3	76.8	70.2	80.5	90.5	76.2	84.0	87.9	81.1
	InfoSieve [38]	94.8	<u>97.7</u>	93.4	78.3	82.2	70.5	80.5	93.8	73.8	84.5	91.2	79.2
	SimGCD [53]	97.1	95.1	98.1	80.1	81.2	77.8	83.0	93.1	77.9	86.7	89.8	84.6
	+ ConGCD	97.3	95.1	98.4	81.3	82.5	78.9	83.5	92.8	78.6	87.4	90.1	85.3
	LegoGCD [6]	97.1	94.3	<u>98.5</u>	81.8	81.4	82.5	86.3	<u>94.5</u>	<u>82.1</u>	<u>88.4</u>	90.1	87.7
	+ ConGCD	97.2	94.8	97.9	<u>82.2</u>	82.1	<u>82.3</u>	85.8	94.4	81.7	<u>88.4</u>	90.4	<u>87.3</u>
	SPTNet [52]	97.3	95.0	98.6	81.3	84.3	75.6	85.4	93.2	81.4	88.0	90.8	85.2
	+ ConGCD	97.4	95.2	<u>98.5</u>	82.5	85.9	77.3	<u>85.9</u>	93.4	82.5	88.6	91.5	86.1
	SelEx [39]	94.1	<u>97.7</u>	92.2	80.0	84.8	70.4	82.3	93.9	76.5	85.4	92.1	79.7
	+ ConGCD	95.7	97.6	94.8	<u>80.5</u>	<u>84.9</u>	71.7	83.5	94.2	78.0	86.6	<u>92.2</u>	81.5
	Avg. Δ	+0.50	+0.15	+0.55	+0.83	+0.93	+0.98	+0.43	-0.03	+0.73	+0.63	+0.35	+0.75

Table 2. Results on the **coarse-grained** classification benchmark. ConGCD offers excellent cross-scheme and cross-model compatibility.

5.1. Setup

Benchmarks. ConGCD is evaluated on a total of seven image recognition benchmarks, namely four fine-grained datasets, CUB-200-2011 [51], Stanford Cars [23], FGVC Aircraft [32], Herbarium19 [43], and three coarse-grained datasets, CIFAR10, CIFAR100 [26] and ImageNet100 [17]. When dealing with CUB, Stanford Cars, and FGVC Aircraft, to segregate target classes into sets of known and unknown, we follow the splits defined by the Semantic Shift Benchmark (SSB) [47]. For the remaining datasets, the splits from the previous study [48] are employed. Under the CIFAR100 benchmark, we designate 80% of the classes as known. For the other benchmarks, the proportion of known classes is 50%. Our labeled set, known as \mathcal{D}_l , consists of 50% images from the known classes for all benchmarks.

Evaluation Protocols. Following existing research [39], we use Balanced Semi-Supervised K-means [50] to cluster the entire image collection \mathcal{D} and calculate the model’s recognition accuracy on the dataset \mathcal{D}_u , which does not have access to ground truth labels. The Hungarian algorithm [54] is used to obtain the optimal mapping relationship between the unlabeled cluster set and the ground truth labels. Instances in the unlabeled set \mathcal{D}_u are divided into **Known** and **Novel** categories based on their category primitives for clustering accuracy evaluation. In addition, the accuracy of **All** classes is reported, as it represents the comprehensive performance of the model.

Implementation Details. We employ the pre-trained DINOv1 [7] on ImageNet-1K [27] as image encoder, and keep it frozen. Features from the 10-th block are utilized for the

decoupling of visual primitives. We set number of visual primitives to 8. The Slot-Attention is trained with the Adam optimizer for 400 epochs with learning rate of 0.0005, batch size of 64. We construct the ConGCD for the last two blocks, following with the previous research [39, 48]. The $\lambda = 0.35$, label smoothing coefficient α set to 0.5 in the fine-grained dataset and 0.1 in the coarse-grained dataset. Additionally, we conduct the above training on DINOv2 [33] pre-trained on ImageNet-22K [40]. All experiments are conducted on RTX-4090, and random number seeds are fixed at each stage to avoid data label leaks.

5.2. Baselines

In addition to the contrast-based method CMS [10] and the use of prototype classifiers in SimGCD [53], we also compare ConGCD with several other state-of-the-art existing methods. LegoGCD [6] [48] leverages a modular approach to discover novel categories in unlabelled datasets by combining different components. SPTNet [52] iteratively optimizes both model parameters and data parameters through spatial prompt tuning. SelEx [39] addresses GCD through self-expertise strategies for hierarchical pseudo-label.

5.3. Results

Quantitative Results. ConGCD demonstrates significant improvements in both fine-grained and coarse-grained classification tasks, particularly excelling in classifying new categories while maintaining performance on known categories. In fine-grained tasks (Tab 1), ConGCD enhances overall accuracy by an average of 1.18% across four baselines, with accuracy on new categories increasing by 1.98%.




Figure 4. Comparison between (a) $\log(\text{rank}(\mathcal{A}))$ and (b) $\hat{H}(\mathcal{A})$. The count of the largest eigenvalues necessary to account for 99% of the total eigenvalue energy serves as a surrogate for the rank.

For coarse-grained tasks (Tab 2), ConGCD improves accuracy on new categories by 0.98% on CIFAR-100, while accuracy on known categories remains stable or shows slight improvement. These results highlight ConGCD’s ability to effectively classify new categories without compromising performance on known ones, demonstrating its effectiveness as a plug-and-play method with strong compatibility across diverse approaches.




Figure 5. Visualization of visual primitives. In a dataset, primitives are bound to specific semantics in an unsupervised manner.

Qualitative Results. (1) Contextual consensus units capture inherent distributions, while dominant consensus units focus on class-related variables. As shown in Fig. 3, weak neurons activate in low-information areas like the background, whereas dominant neurons target high-semantic difference areas (*e.g.*, wheels, headlights) to enhance GCD. (2) To highlight ConGCD’s performance improvement, attention maps in Fig. 8 reveal that ConGCD emphasizes the main body of the object (foreground), unlike SelEx, which focuses more on the background and specific minor areas. This enables a more uniform extraction of task-related variables. (3) As shown in Fig. 5, it can be observed that within the same dataset, visual primitives possess a high degree of cross-sample consistency. A primitive focuses on a region that has similar semantics.

5.4. Analysis

Visual primitives increase von Neumann entropy. We analyze the sources of the effectiveness of ConGCD from the perspectives of representation learning and information quantification. We employ von Neumann entropy [35] as a measure of the amount of information in the representation. For the class token \mathbf{z} used in GCD, its autocorrelation




Figure 6. Visualisation of the embedding space with t-SNE [45].

matrix $\mathcal{A} \triangleq \sum_{i=1}^N \frac{1}{N} \mathbf{z}_i \mathbf{z}_i^\top$. The von Neumann entropy (VNA) is expressed as $\hat{H}(\mathcal{A}) \triangleq -\sum_j \lambda_j \log \lambda_j$, where λ is the eigenvalues of \mathcal{A} . As shown in Fig. 4d and Fig. 4a, compared with SelEx, the representation of ConGCD has a higher $\hat{H}(\mathcal{A})$ and a higher $\log(\text{rank}(\mathcal{A}))$, which means a larger amount of information. This indicates that ConGCD enables the GCD process to more discriminatively separate clusters by providing more discriminative features.

The versatility of visual primitives. To verify the universality of ConGCD as a plug-and-play solution, we integrated ConGCD with existing GCD schemes, and the results are reported in Tab. 1 and Tab. 2. It is evident that ConGCD provides them with a consistent improvement in accuracy, which can be attributed to the unstructured decoupling of visual primitives from objects, enabling the model to possess a certain ability to model topological relationships. The model no longer solely relies on image-level embeddings from a single perspective and applies this category-independent topological relationship to the discovery of novel categories.

ConGCD provides accurate distribution estimation. The core of GCD is to estimate category distributions as accurately as possible and separate different categories. We evaluate ConGCD’s contribution to distribution estimation from two perspectives: clustering visualization and category number estimation. (1) We visualize embeddings using t-SNE [45], as shown in Fig. 6. Notably, ConGCD significantly enhances clustering quality without modifying the optimization objective. It reduces within-class distances, clarifying cluster boundaries. (2) As shown in ??,

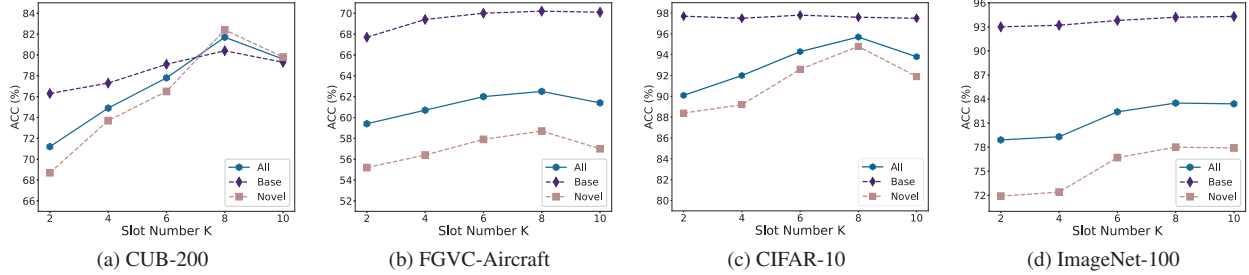


Figure 7. The ablations of the number of visual primitives K on (a, b) fine- and (c, d) coarse-grained benchmarks.

Method	CIFAR-100		ImageNet-100		CUB-200		Stanford-Cars	
	$ \mathcal{Y}_u $	Err(%)	$ \mathcal{Y}_u $	Err(%)	$ \mathcal{Y}_u $	Err(%)	$ \mathcal{Y}_u $	Err(%)
Ground truth	100	-	100	-	200	-	196	-
GCD [48]	100	0	109	9	231	15.5	230	17.3
DCCL [36]	146	46	129	29	172	9	192	0.02
PIM [9]	95	5	102	2	227	13.5	169	13.8
GPC [57]	100	0	103	3	212	6	201	0.03
CMS [10]	94	6	98	2	176	12	149	23.9
ConGCD	97	3	99	1	186	7	160	18.3

Table 3. Estimated number and error rate of $|\mathcal{Y}_u|$.



Figure 8. Visualization on attention maps of SoTA and ConGCD.

ConGCD combined with CMS [10] achieves the most accurate category number estimation $|\mathcal{Y}_u|$ with the lowest error rate, demonstrating its superior ability to model data and partition the feature space, crucial for reaching GCD’s upper bound.

Components	CUB-200			FGVC-Aircraft			CIFAR-10			ImageNet-100		
	All	Known	Novel	All	Known	Novel	All	Known	Novel	All	Known	Novel
conventional MoE	78.2	79.1	77.3	59.4	78.1	53.6	93.6	97.1	91.8	80.4	92.5	74.9
Ours w/o Dominant Consensus Units	81.1	80.3	81.5	62.1	69.7	58.3	94.9	97.8	94.1	82.9	94.0	77.2
Ours w/o Contextual Consensus Units	80.4	79.7	81.0	61.3	70.6	57.1	94.2	97.6	92.3	81.5	95.5	75.4
Ours w/o Consensus Scheduler	79.8	79.5	80.0	60.9	70.1	56.7	94.4	97.7	92.7	81.4	92.7	77.0

Table 4. Ablations on components of multiplex consensus.

Ablations on the number of visual primitives. As a plug-and-play solution, ConGCD requires no hyperparameter tuning, which is an important prerequisite for its wide application. The only parameter that may potentially affect the performance is the number K of visual primitives. We conducted ablation studies on it for both fine- and coarse-grained datasets, as shown in Fig. 7. It can be seen that in the open world, an appropriate K is sufficient to handle

multiple downstream tasks. When K is in the range of 7 to 10, ConGCD achieves consistent and stable performance improvements, and K is set to 8 for all dataset.

Ablation Study on the Multiplex Consensus Mechanism.

We conducted an ablation study on the key component of ConGCD in Sec. 4.2, with results summarized in Tab. 4. The mixture of experts (MoE) is the most comparable approach to our multiplex consensus mechanism, as it directly integrates with visual primitives. However, conventional MoE proves detrimental to the GCD task, as it only accepts image-level embeddings and overfits to labeled data from known classes. The experts merely memorize features of known classes, failing to generalize to novel categories. In contrast, our multiplex consensus mechanism offers three key advantages: **(1) Dominant consensus units** process diverse visual primitives, enabling primitive-level perception and improving classification accuracy for novel categories. **(2) Contextual consensus units** capture intrinsic distributions shared across known and unknown classes, ensuring balanced clustering and mitigating bias. **(3) Consensus scheduler** introduces a learnable threshold to refine the distinction between weak and dominant neurons, optimizing unit selection and weight distribution to enhance accuracy.

6. Conclusion

Current GCD approaches treat images as atomic entities, ignoring the human capacity for compositional reasoning. We propose ConGCD, a cognition-inspired paradigm bridging this gap with two key innovations: self-deconstruction decomposes objects into competing visual primitives via semantic reconstruction, while multiplex consensus dynamically integrates dual pathways, dominant consensus units capturing class-specific attributes and contextual consensus units modeling distributional invariants. ConGCD achieves competitive performance on coarse- and fine-grained benchmarks, enhancing interpretability and generalization in open-world recognition.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grant 82172033, Grant U19B2031, Grant 61971369, Grant 52105126, Grant 82272071, and Grant 62271430; and in part by the Dreams Foundation of Jianghuai Advance Technology Center; and in part by the Open Fund of the National Key Laboratory of Infrared Detection Technologies.

References

- [1] Patricia A Alexander. Relational thinking and relational reasoning: harnessing the power of patterning. *NPJ science of learning*, 1(1):1–7, 2016. 3
- [2] Görkay Aydemir, Weidi Xie, and Fatma Guney. Self-supervised object-centric learning for videos. *Advances in Neural Information Processing Systems*, 36:32879–32899, 2023. 2
- [3] Anwesha Banerjee, Liyana Sahir Kallooriyakath, and Soma Biswas. Amend: Adaptive margin and expanded neighborhood for efficient generalized category discovery. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2101–2110, 2024. 5
- [4] Mario Belloni and Richard W Robinett. The infinite well and dirac delta function potentials as pedagogical, mathematical and physical models in quantum mechanics. *Physics Reports*, 540(2):25–122, 2014. 3
- [5] Kaidi Cao, Maria Brbic, and Jure Leskovec. Open-world semi-supervised learning. 2021. 6
- [6] Xinzi Cao, Xiawu Zheng, Guan hong Wang, Weijiang Yu, Yunhang Shen, Ke Li, Yutong Lu, and Yonghong Tian. Solving the catastrophic forgetting problem in generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16880–16889, 2024. 5, 6
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 5, 6
- [8] Lin Chen. Topological structure in visual perception. *Science*, 218(4573):699–700, 1982. 1, 2, 3
- [9] Florent Chiaroni, Jose Dolz, Ziko Imtiaz Masud, Amar Mitiche, and Ismail Ben Ayed. Parametric information maximization for generalized category discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1729–1739, 2023. 6, 8
- [10] Sua Choi, Dahyun Kang, and Minsu Cho. Contrastive mean-shift learning for generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23094–23104, 2024. 1, 5, 6, 8
- [11] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014. 3
- [12] Rahul Dey and Fathi M Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600. IEEE, 2017. 3
- [13] Oxford English Dictionary. Oxford english dictionary. *Simpson, Ja & Weiner, Esc*, 3, 1989. 2
- [14] Gamaleldin Elsayed, Aravindh Mahendran, Sjoerd Van Steenkiste, Klaus Greff, Michael C Mozer, and Thomas Kipf. Savi++: Towards end-to-end object-centric learning from real-world videos. *Advances in Neural Information Processing Systems*, 35:28940–28954, 2022. 2
- [15] Yixin Fei, Zhongkai Zhao, Siwei Yang, and Bingchen Zhao. Xcon: Learning with experts for fine-grained category discovery. *arXiv preprint arXiv:2208.01898*, 2022. 5, 6
- [16] Vittorio Ferrari and Andrew Zisserman. Learning visual attributes. *Advances in neural information processing systems*, 20, 2007. 2
- [17] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018. 6
- [18] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*, 2019. 2, 4
- [19] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. 2
- [20] Nathaniel R Goodman. Statistical analysis based on a certain multivariate complex gaussian distribution (an introduction). *The Annals of mathematical statistics*, 34(1):152–177, 1963. 3
- [21] Jinwoo Kim, Janghyuk Choi, Ho-Jin Choi, and Seon Joo Kim. Shepherding slots to objects: Towards stable and robust object-centric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19198–19207, 2023. 2
- [22] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *International conference on machine learning*, pages 5338–5348. PMLR, 2020. 1
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. 6
- [24] Daniel C Krawczyk. The cognition and neuroscience of relational reasoning. *Brain research*, 1428:13–23, 2012. 2
- [25] Daniel C Krawczyk, M Michelle McClelland, and Colin M Donovan. A hierarchy for relational reasoning in the prefrontal cortex. *Cortex*, 47(5):588–597, 2011. 2
- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6

- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 6
- [28] Deng Li, Aming Wu, Yaowei Wang, and Yahong Han. Prompt-driven dynamic object-centric learning for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17606–17615, 2024. 2
- [29] Tianqin Li, Ziqi Wen, Yangfan Li, and Tai Sing Lee. Emergence of shape bias in convolutional neural networks through activation sparsity. *Advances in Neural Information Processing Systems*, 36, 2024. 4
- [30] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in neural information processing systems*, 33:11525–11538, 2020. 2, 3
- [31] Shijie Ma, Fei Zhu, Zhun Zhong, Xu-Yao Zhang, and Cheng-Lin Liu. Active generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16890–16900, 2024. 1, 2
- [32] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 6
- [33] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 6
- [34] Jona Otholt, Christoph Meinel, and Haojin Yang. Guided cluster aggregation: A hierarchical approach to generalized category discovery. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2618–2627, 2024. 2
- [35] Dénes Petz. Entropy, von neumann and the von neumann entropy: Dedicated to the memory of alfred wehrl. In *John von Neumann and the foundations of quantum physics*, pages 83–96. Springer, 2001. 7
- [36] Nan Pu, Zhun Zhong, and Nicu Sebe. Dynamic conceptional contrastive learning for generalized category discovery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7579–7588, 2023. 3, 6, 8
- [37] Nan Pu, Wenjing Li, Xingyuan Ji, Yalan Qin, Nicu Sebe, and Zhun Zhong. Federated generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28741–28750, 2024. 2
- [38] Sarah Rastegar, Hazel Doughty, and Cees Snoek. Learn to categorize or categorize to learn? self-coding for generalized category discovery. *Advances in Neural Information Processing Systems*, 36, 2024. 5, 6
- [39] Sarah Rastegar, Mohammadreza Salehi, Yuki M Asano, Hazel Doughty, and Cees GM Snoek. Selex: Self-expertise in fine-grained generalized category discovery. *arXiv preprint arXiv:2408.14371*, 2024. 1, 5, 6
- [40] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021. 6
- [41] Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, Thomas Brox, et al. Bridging the gap to real-world object-centric learning. *arXiv preprint arXiv:2209.14860*, 2022. 2
- [42] Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, Thomas Brox, and Francesco Locatello. Bridging the gap to real-world object-centric learning. In *The Eleventh International Conference on Learning Representations*, 2023. 4
- [43] Kiat Chuan Tan, Yulong Liu, Barbara Ambrose, Melissa Tulig, and Serge Belongie. The herbarium challenge 2019 dataset. *arXiv preprint arXiv:1906.05372*, 2019. 6
- [44] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 1
- [45] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9 (86):2579–2605, 2008. 7
- [46] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 3, 5
- [47] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: A good closed-set classifier is all you need? 2021. 6
- [48] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7492–7501, 2022. 1, 2, 3, 5, 6, 8
- [49] Sagar Vaze, Andrea Vedaldi, and Andrew Zisserman. No representation rules them all in category discovery. *Advances in Neural Information Processing Systems*, 36, 2024. 5
- [50] Viet-Vu Vu, Nicolas Labroche, and Bernadette Bouchon-Meunier. Active learning for semi-supervised k-means clustering. In *2010 22nd IEEE international conference on tools with artificial intelligence*, pages 12–15. IEEE, 2010. 6
- [51] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 6
- [52] Hongjun Wang, Sagar Vaze, and Kai Han. Sptnet: An efficient alternative framework for generalized category discovery with spatial prompt tuning. *arXiv preprint arXiv:2403.13684*, 2024. 1, 5, 6
- [53] Xin Wen, Bingchen Zhao, and Xiaojuan Qi. Parametric classification for generalized category discovery: A baseline study. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16590–16600, 2023. 1, 2, 5, 6
- [54] Mike B Wright. Speeding up the hungarian algorithm. *Computers & Operations Research*, 17(1):95–96, 1990. 6
- [55] Mateo Espinosa Zarlenga, Pietro Barbiero, Gabriele Ciravegna, Giuseppe Marra, Francesco Giannini, Michelangelo Diligenti, Frederic Precioso, Stefano Melacci, Adrian Weller, Pietro Lio, et al. Concept embedding models. In

NeurIPS 2022-36th Conference on Neural Information Processing Systems, 2022. [1](#)

- [56] Sheng Zhang, Salman Khan, Zhiqiang Shen, Muzammal Naseer, Guangyi Chen, and Fahad Shahbaz Khan. Promptcal: Contrastive affinity learning via auxiliary prompts for generalized novel category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3479–3488, 2023. [5](#), [6](#)
- [57] Bingchen Zhao, Xin Wen, and Kai Han. Learning semi-supervised gaussian mixture models for generalized category discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16623–16633, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [8](#)