

When Domain Generalization meets Generalized Category Discovery: An Adaptive Task-Arithmetic Driven Approach

Vaibhav Rathore¹ Shubhranil B¹ Saikat Dutta^{1,2*} Sarthak Mehrotra^{1*} Zsolt Kira³
Biplab Banerjee¹
¹IIT Bombay ²IITB-Monash Research Academy ³Georgia Institute of Technology

Abstract

Generalized Class Discovery (GCD) clusters base and novel classes in a target domain, using supervision from a source domain with only base classes. Current methods often falter with distribution shifts and typically require access to target data during training, which can sometimes be impractical. To address this issue, we introduce the novel paradigm of Domain Generalization in GCD (DG-GCD), where only source data is available for training, while the target domain—with a distinct data distribution—remains unseen until inference. To this end, our solution, DG²CD-Net, aims to construct a domain-independent, discriminative embedding space for GCD. The core innovation is an episodic training strategy that enhances cross-domain generalization by adapting a base model on tasks derived from source and synthetic domains generated by a foundation model. Each episode focuses on a cross-domain GCD task, diversifying task setups over episodes and combining open-set domain adaptation with a novel margin loss and representation learning for optimizing the feature space progressively. To capture the effects of fine-tunings on the base model, we extend task arithmetic by adaptively weighting the local task vectors concerning the fine-tuned models based on their GCD performance on a validation distribution. This episodic update mechanism boosts the adaptability of the base model to unseen targets. Experiments across three datasets confirm that DG²CD-Net outperforms existing GCD methods customized for DG-GCD.

1. Introduction

Recent advances in deep learning [8, 52] have improved visual inference models; however, challenges persist in limited supervision, domain shifts, and detecting novel classes in inference. Semi-supervised and self-supervised methods [22, 48, 69] address limited supervision but assume a closed-set environment. Domain adaptation (DA) and do-



Figure 1. We present a novel variant of GCD, **Domain Generalization for Generalized Category Discovery**, where a model is trained on a **source domain** (photo) and evaluated on a **target domain** (cartoon). During inference, the model must cluster **seen classes** while also identifying and distinguishing **novel classes**.

main generalization (DG) [55, 79] aim to create domain-agnostic embeddings for distribution generalization. However, in open-set applications [3, 41], they typically classify novel samples in a single outlier class.

Applications like autonomous driving and healthcare require finer categorization of novel samples. Novel Class Discovery (NCD) [15, 19] addresses this by clustering novel classes in an unlabeled target domain using labeled source data with known classes. Generalized Class Discovery (GCD) [38, 49, 72] extends this to both known and novel classes in the target domain, with the support of source domain knowledge. Cross-Domain GCD (CD-GCD) [40] further advances this by handling cases where the source and target domains differ in distribution. Nonetheless, all these settings assume simultaneous access to source and target data to carry out training, which is often impractical. For example, an autonomous vehicle [45, 70] trained on California’s roads in summer may face unexpected challenges on snowy Moscow streets. Similarly, medical AI models [2, 68] trained on high-quality images may need to analyze lower-resolution scans from rural areas to detect new pathologies. In these cases, the target domains are majorly inaccessible until inference. Thus, robust AI systems must adapt to varying environments, organizing both known and novel data categories without prior exposure and ensuring reliable performance in unexpected scenarios.

To meet these challenges, we propose DG-GCD (Table 1), a novel and practical setting that leverages a labeled source domain to train a feature extractor on known classes.

*These authors contributed equally to this work.

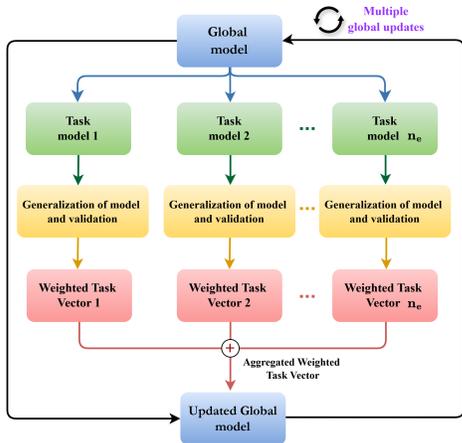


Figure 2. **Proposed episodic training:** A pre-trained global model is updated using task vectors from n_e episode-specific fine-tuned models, leveraging a novel dynamic weighting scheme. This scheme adjusts the task vectors based on their GCD generalization performance on a held-out unseen validation distribution.

The task is to leverage the trained model to cluster unlabeled samples from a previously unseen, visually distinct target domain containing both known and novel classes, with an undefined known-novel class ratio—suggesting *solving the GCD problem within a DG framework with uncharted domain and semantic shifts between training and testing*.

Existing GCD and CD-GCD methods perform sub-optimally in DG-GCD, when we train them solely on the source domain without considering the target domain as per DG requirements (Table 1). Besides, single-source open-set DG methods [3], which generalize from a labeled source to a mixed-class target domain, categorize all novel samples as a single outlier, diverging from DG-GCD’s objectives.

Our proposal: As a remedy, we present DG-GCD-Net (*aka* DG²CD-Net). Generalizing from a single-source domain is inherently challenging due to limited data diversity, the need to adapt to unseen novel classes in the target domain, and the varying known-novel class splits that complicate learning. Standard approaches often struggle to effectively handle both domain shifts and class-level variations.

To address this, we propose a novel **progressive episodic training strategy** (Fig. 2) designed to enhance a pre-trained global feature encoder. In each episode, the CD-GCD task is tackled locally by using a subset of the source domain with a reduced class set. Additionally, a stylistic variant of the source domain, generated via a diffusion model [5], acts as a pseudo-target domain containing all known classes, but is treated unlabeled. These episodic splits are dynamically generated to simulate diverse conditions.

For CD-GCD, known classes from the source (labeled) and synthesized (unlabeled) domains are aligned using an open-set domain adaptation (OSDA) objective, while contrastive learning boosts feature discriminativeness. Exist-

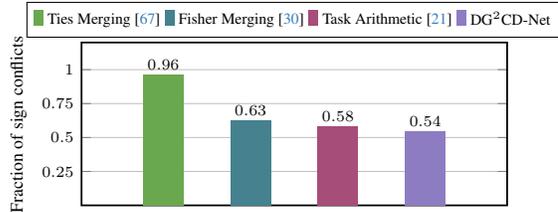


Figure 3. **Sign Conflicts**—opposing gradient updates leading to greater parameter divergence across episodes in various **model merging methods** for our episodic training on PACS. Lower values indicate improved training stability and accuracy.

ing OSDA methods like [41] often fail with fine-grained novel classes, causing misalignment with known classes. To address this, we propose a **novel margin objective** to enforce clearer separation between episode-specific known and novel classes, improving novel class discovery.

To aggregate task-specific knowledge from multiple fine-tuned models, we apply task arithmetic (TA) [21], combining task vectors that capture weight shifts between pre-trained and fine-tuned models to improve adaptability. While certain TA variants focus only on top-ranked parameters [67], this approach leads to instability in our context of multiple parallel episodes with varied data distributions. To maintain consistency, we use TA across the entire weight space. To enhance generalization beyond traditional TA, we introduce a novel **dynamically weighted TA** approach tailored for DG-GCD, which deals with scaling task vectors according to each fine-tuned model’s generalization performance, assessed through clustering on a diverse validation set. This prioritizes broad generalization, preventing over-specialization and supporting robust adaptation across novel domains. In contrast to approaches like [30], which rely on the Fisher matrix for weighted model merging, our approach better accommodates divergent data distributions, making it more suited to our training scheme (Fig. 3, Table 3). Our major contributions are summarized as follows:

- [-] We introduce a novel variant of GCD, DG-GCD, where the target domain, distinct from the source in distribution, is only encountered at evaluation.
- [-] We introduce an innovative episodic training strategy that utilizes adaptive task vector modeling, dynamically updating the pre-trained global model by leveraging the generalizability of fine-tuned models tailored for CD-GCD.
- [-] We perform extensive experiments on three datasets, benchmarking our model against several baselines and conducting thorough ablation studies.

2. Related Works

Domain generalization: DG [54, 79] trains models on one or more source domains to generalize to unseen target domains, covering scenarios like closed-set, open-set, multi-source, and single-source. Multi-source DG enhances

model robustness using techniques like invariant feature learning and meta-learning [14, 26, 27, 31, 77, 78], while single-source DG [11, 37, 59, 64] relies on data augmentation and regularization due to limited domain diversity.

Closed-set DG assumes all target domain classes are seen during training, whereas open-set DG [3, 4, 35, 42, 56, 57, 80] introduces unseen target classes, increasing complexity. Recent methods, such as [44], use prompt learning with CLIP to address open-set DG. However, open-set DG groups all novel samples into a single class, unlike our DG-GCD framework, which aims for fine-grained discovery and clustering of novel classes.

(Cross-domain) class discovery: While both NCD and GCD are well-studied problems [9, 16, 17, 19, 38, 61, 72, 75, 76] (see Section B for more details), class discovery under domain shift remains relatively new. Emerging methods such as CDAD-Net [40], Exclusive Style Removal [58], CROW [60], and HiLo [53] concentrate on aligning known classes across domains and strengthening the generalization of novel class discovery amid distributional shifts. These approaches, assuming target domain presence during training, serve as upper bounds for DG²CD-Net’s performance.

Model aggregation: Fine-tuned models initialized from the same pre-trained model often follow similar optimization paths, making them suitable for merging [21, 23, 62, 63]. Model merging has been applied across various fields, including out-of-domain generalization [1, 7, 21, 23], multitask learning [23], federated learning [28], multimodal merging [46], and continual learning [65, 66].

Task Arithmetic [21] improves merging by using task vectors—differences between the weights of a pre-trained model and its fine-tuned versions—to update the base model. By integrating these task vectors, TA aims to boost the pre-trained model’s performance on new tasks while preserving stability. Further refinements to TA have been proposed, such as Ortiz-Jimenez *et al.* [34], using a weight disentanglement framework in tangent space, and Yadav *et al.* [67], addressing parameter interference during merging. DG²CD-Net extends on the conventional TA by introducing an episodic training scheme to handle dynamic domain and semantic shifts, which static updates in [21, 67] fail to address. We also adaptively weigh task vectors based on the GCD generalizability of the fine-tuned models on held-out and disjoint data distributions, ensuring the global model avoids bias towards subpar local models.

3. Problem Definition & Methodology

In this section, we introduce the DG-GCD problem, where a model trained solely on labeled data from a source domain \mathcal{S} must identify and cluster both known and novel categories in an unlabeled target domain \mathcal{T} during testing, without prior exposure, despite domain shifts ($\mathcal{P}(\mathcal{S}) \neq \mathcal{P}(\mathcal{T})$). The source domain dataset is $\mathcal{D}_{\mathcal{S}} = \{(x_i^s, y_i^s)\}_{i=1}^{n_s}$, where

$x_i^s \in \mathcal{X}_s$ are inputs and $y_i^s \in \mathcal{Y}_s$ are labels of known categories. The target domain dataset $\mathcal{D}_{\mathcal{T}} = \{x_j^t\}_{j=1}^{n_t}$ contains both known and novel categories, denoted as $\mathcal{Y}_t^{old} = \mathcal{Y}_s$ and $\mathcal{Y}_t^{new} = \mathcal{Y}_t \setminus \mathcal{Y}_s$, respectively. The core challenge is to develop a domain-generalizable and discriminative embedding space suitable for on-the-fly GCD evaluation.

3.1. The DG²CD-NET model

We introduce DG²CD-Net, initialized from a pre-trained encoder architecture \mathcal{F}^0 with parameters θ_{global} based on the Vision Transformer (ViT) [12]. To enhance the generalization capabilities of θ_{global} for GCD tasks, we propose an adaptive task-vector aggregation strategy within an episodic training framework. This framework leverages both source data $\mathcal{D}_{\mathcal{S}}$ and a set of \mathcal{N} synthetic pseudo-domains \mathcal{D}_{syn} , generated by a pre-trained foundation model to systematically vary style elements while preserving the semantic structure of $\mathcal{D}_{\mathcal{S}}$. Our training scheme with such varied data aims to progressively strengthen the domain independence and discriminative capacity of θ_{global} . Further details on the domain synthesis strategy are provided in Section 3.1.1.

Our training protocol comprises multiple rounds of global updates for θ_{global} , each consisting of several local episodes. In each episode e within a global update g , we select a labeled subset $\mathcal{D}_{\mathcal{S}}^{e,g} \subset \mathcal{D}_{\mathcal{S}}$ and a randomly chosen synthetic domain $\mathcal{D}_{\text{syn}}^{e,g}$, enabling the model to adapt locally to the CD-GCD task [40] and thereby account for distribution shifts in a GCD context. In this setup, the subset $\mathcal{D}_{\mathcal{S}}^{e,g}$ contains a reduced set of classes, $\mathcal{Y}_s^{e,g} \subset \mathcal{Y}_s$ and is treated labeled, while $\mathcal{D}_{\text{syn}}^{e,g}$ includes the full class set \mathcal{Y}_s but is assumed unlabeled. $\mathcal{Y}_s \setminus \mathcal{Y}_s^{e,g}$ defines the episode-specific novel classes in $\mathcal{D}_{\text{syn}}^{e,g}$. Varying these subsets over episodes exposes the model to diverse domain and semantic shifts, progressively aiding in enhancing the robustness of the updated θ_{global} for DG-GCD.

To address the CD-GCD problem, we apply task-specific objectives to each episode pair $(\mathcal{D}_{\mathcal{S}}^{e,g}, \mathcal{D}_{\text{syn}}^{e,g})$ initialized from the current global model, $\theta_{\text{global}}^{g-1}$, yielding an episode-specific parameterization $\theta_{\text{local}}^{e,g}$. The generalization performance of $\theta_{\text{local}}^{e,g}$ is evaluated on a validation set, $\mathcal{D}_{\text{valid}}$, constructed from diverse synthetic domains distinct from those used in the episodes and held constant throughout training. Keeping $\mathcal{D}_{\text{valid}}$ a mixed distribution and separate helps assess the DG-GCD ability for the local models more effectively in a challenging scenario (see Table 3 for ablation). We assess performance on $\mathcal{D}_{\text{valid}}$ using K-means clustering with Hungarian assignment [24]. The resulting validation metrics guide the adaptive weighting of task vectors for $\theta_{\text{local}}^{e,g}$, and the knowledge from all episodes within the global update g is aggregated to refine $\theta_{\text{global}}^{g-1}$ into θ_{global}^g (or \mathcal{F}^g). The complete training process, together with the CD-GCD losses, are detailed in Section 3.1.2.

3.1.1. Synthetic domain generation

Our aim is to generate diverse synthetic domains, \mathcal{D}_{syn} , to introduce distributional variety that enhances generalization beyond the source domain, $\mathcal{D}_{\mathcal{S}}$. Manual ad-hoc augmentations, however, are limited by the need for domain-specific expertise and lack the flexibility to introduce realistic distribution shifts. To address this, we leverage Instruct-Pix2Pix [5], a pre-trained conditioned diffusion model, which enables controlled and varied transformations via targeted prompts. This approach allows for modifications in texture, lighting, and environmental factors while preserving semantic integrity, resulting in realistic and adaptable distribution shifts (see Table 3 for comparisons). Moving forward, we start with curating a number of domain descriptors that can introduce such variations in $\mathcal{D}_{\mathcal{S}}$ by querying ChatGPT-4o [32] as follows:

Generate a list of prompts in the format “Add [description] scene to the image.” Replace ‘[description]’ with single-word terms or short phrases representing colors, weather, time of day, or location. Ensure variety with minimal adjectives. Each should evoke a distinct visual trait.

This generates domain primitives such as Snowy-Scene and Rainy-Weather for textures and occlusions, Night-Time for low visibility, and Beach, Forest, and Urban for environmental diversity. These curated prompts are further used in Instruct-Pix2Pix to synthesize the images. Importantly, this approach avoids leaking any target-domain-specific knowledge, preserving DG integrity. As shown in Fig. 4, these synthesized domains enhance the distribution diversity of $\mathcal{D}_{\mathcal{S}} \cup \mathcal{D}_{\text{syn}}$, more effectively surrogating the original target counterparts than manual ad-hoc image augmentations (see Sections E and F for further quantitative and qualitative analyses).

3.1.2. Proposed episodic training strategy using task vectors to update θ_{global} and the CD-GCD objectives

Given datasets $\mathcal{D}_{\mathcal{S}}$ and \mathcal{D}_{syn} , our goal is to optimize θ_{global} to construct a domain-generalized embedding space that maintains robust class discriminability across varied known-novel class ratios. Traditional Task Arithmetic [21] updates θ_{global} by integrating task vectors from fine-tuned models; however, applying TA directly for CD-GCD—using static task vectors derived from joint fine-tuning on $\mathcal{D}_{\mathcal{S}}$ and \mathcal{D}_{syn} —proves inadequate for handling dynamic distribution shifts and evolving semantics (Table 3). Moreover, TA’s absence of validation-driven assessment increases the risk of overfitting to familiar domains, leading to diminished performance on unseen distributions, as reflected in DG²CD-Net (TA) results in Table 1.

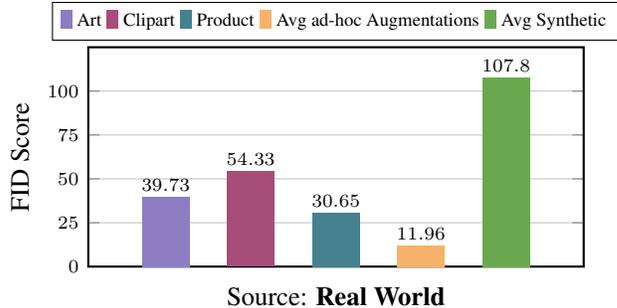


Figure 4. The **FID** scores [13] between the *Real-World* and the three other Office-Home [50] domains, along with the average FID between *Real-World* and the InstructPix2Pix-driven and manually-crafted synthesized domains, confirm that the InstructPix2Pix-driven synthetic domains introduce diverse distribution shifts into our episodic training process, leading to enhanced generalizability of θ_{global} through the proposed training scheme.

We introduce an episodic training framework with a novel dynamically weighted task-vector aggregation strategy to address these limitations (Fig. 5). Training is structured across n_g global updates, each containing n_e episodes. At the start of the e^{th} episode, the task model is initialized with parameters $\theta_{\text{global}}^{g-1}$ from the previous global iteration. The model is then fine-tuned on the episode’s dataset pair $(\mathcal{D}_{\mathcal{S}}^{e_g}, \mathcal{D}_{\text{syn}}^{e_g})$ under the CD-GCD objective to obtain $\theta_{\text{local}}^{e_g}$. Subsequently, the task vector for the fine-tuned model, δ_g^e , is computed by the parameter difference:

$$\delta_g^e = \theta_{\text{global}}^{g-1} - \theta_{\text{local}}^{e_g} \quad (1)$$

Rather than directly integrating δ_g^e into $\theta_{\text{global}}^{g-1}$, we propose weighting δ_g^e according to the GCD generalization ability of $\theta_{\text{local}}^{e_g}$, assessed on $\mathcal{D}_{\text{valid}}$. The All performance metric [49], which captures performance on both known and novel categories of $\mathcal{D}_{\text{valid}}$, serves as an effective measure of generalization; a high All score indicates superior generalizability. We further normalize the All scores across all episodes of the global update index g (Fig. 5). Instead of using naive normalization methods like min-max normalization, we apply a softmax function to the All scores (see Table 3 for a comparative study). The softmax weighting emphasizes high-performing episodes by exponentially scaling weights, amplifying task vectors with stronger generalization potential while attenuating less effective ones. The generalization weight corresponding to $\theta_{\text{local}}^{e_g}$ is then calculated as follows:

$$w_g^e = \frac{\exp(\text{all}_g^e)}{\sum_{e'=1}^{n_e} \exp(\text{all}_g^{e'})} \quad (2)$$

This adaptive weighting then informs the global update:

$$\theta_{\text{global}}^g = \theta_{\text{global}}^{g-1} - \sum_{e=1}^{n_e} w_g^e \delta_g^e \quad (3)$$

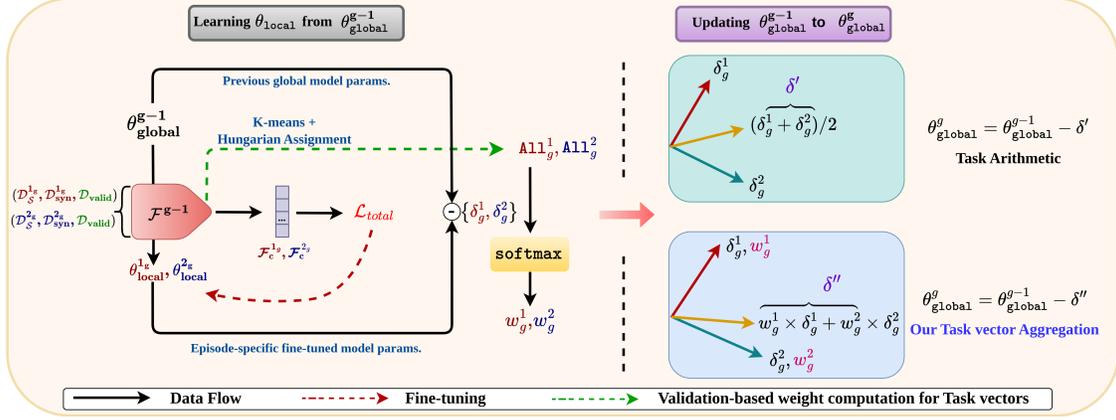


Figure 5. **Transition from $\theta_{\text{global}}^{g-1}$ to θ_{global}^g in our training strategy:** The **left panel** illustrates the two-way episodic training process. Starting with episode-specific datasets $(\mathcal{D}_S^{1g}, \mathcal{D}_{\text{syn}}^{1g})$ and $(\mathcal{D}_S^{2g}, \mathcal{D}_{\text{syn}}^{2g})$, we fine-tune the previous global model \mathcal{F}^{g-1} together with episode-specific adversarial classifiers \mathcal{F}_c^{1g} and \mathcal{F}_c^{2g} on the local CD-GCD tasks. This produces fine-tuned models with updated weights $\theta_{\text{local}}^{1g}$ and $\theta_{\text{local}}^{2g}$. We calculate the task vectors (δ_g^1, δ_g^2) for the fine-tuned models. GCD generalization is subsequently assessed on $\mathcal{D}_{\text{valid}}$ using the All metric, resulting in generalization weights (w_g^1, w_g^2) for the fine-tuned models. The **right panel** shows how the global models are updated through task vector aggregations for baseline TA [21] and ours. **Red** and **Blue** denote the episodes-specific data/processing.

Using a consistent validation set across episodes stabilizes task vector weighting, reduces gradient conflicts, and yields a more robust global model than baseline TA. Despite its resemblance with meta-learning [51], unlike traditional meta-learning, which focuses on rapid task adaptation, our framework centers on robust domain generalization, equipping the model to proactively manage evolving distributions. Further discussion on the *meta-knowledge* acquired through our training process is available in Section I.

The CD-GCD objective within each episode, i.e. e_g^{th} episode: Our CD-GCD objectives in each episode comprise three key loss functions: (i) a supervised contrastive loss for $\mathcal{D}_S^{e_g}$, (ii) an unsupervised contrastive objective for $\mathcal{D}_S^{e_g} \cup \mathcal{D}_{\text{syn}}^{e_g}$, and (iii) open-set domain adaptation coupled with a novel confidence optimization objective for fine-grained known and novel class samples in $\mathcal{D}_S^{e_g}$ and $\mathcal{D}_{\text{syn}}^{e_g}$.

- Supervised contrastive loss: We utilize a supervised contrastive loss that operates in the embedding space of the encoder from the last global update, \mathcal{F}^{g-1} , and aims to enhance the compatibility between an anchor sample $x \in \mathcal{D}_S^{e_g}$ and its corresponding class prototype (center) μ^y obtained by averaging the embeddings of the class-wise samples, where y is the label of x . Simultaneously, the loss reduces the similarity between the anchor and prototypes of all negative classes $\mathcal{Y}_s^{e_g} \setminus y$. Given the temperature hyperparameter τ , $\mathcal{L}_{\text{con}}^s$ is given by:

$$\mathcal{L}_{\text{con}}^s = -\mathbb{E}_{(x,y) \sim \mathcal{P}(\mathcal{D}_S^{e_g})} \log \frac{\exp(\cos(\mathcal{F}^{g-1}(x), \mu^y)/\tau)}{\sum_{y' \in \mathcal{Y}_s^{e_g}} \exp(\cos(\mathcal{F}^{g-1}(x), \mu^{y'})/\tau)} \quad (4)$$

- Unsupervised contrastive loss: We employ an unsupervised contrastive loss to enable discriminative feature

learning across domains. Each image x and its geometrically augmented version x^+ (translation, rotation *etc.*) form a positive pair, while all other images in the batch \mathcal{B} , denoted x^- , serve as negatives. The objective is to maximize similarity between x and x^+ , while minimizing it between x and all negatives x^- . This is formulated as:

$$\mathcal{L}_{\text{con}}^u = \mathbb{E}_{x \sim \mathcal{P}(\mathcal{D}_S^{e_g} \cup \mathcal{D}_{\text{syn}}^{e_g})} - \log \frac{\exp(\cos(\mathcal{F}^{g-1}(x), \mathcal{F}^{g-1}(x^+))/\tau)}{\sum_{x' \in \mathcal{B}} \exp(\cos(\mathcal{F}^{g-1}(x), \mathcal{F}^{g-1}(x'))/\tau)} \quad (5)$$

- Domain-alignment objective with the proposed margin loss: While the above objectives encourage learning discriminative features, effective domain alignment is crucial for transferring supervised knowledge from the labeled dataset $\mathcal{D}_S^{e_g}$ to the unlabeled counterpart $\mathcal{D}_{\text{syn}}^{e_g}$. Given the label discrepancy between $\mathcal{D}_S^{e_g}$ and $\mathcal{D}_{\text{syn}}^{e_g}$, we employ an open-set domain adaptation approach to avoid interference from non-overlapping classes of both the datasets. In contrast to the related literature, CDAD-Net [40], which uses a closed-set adversarial loss for aligning the source and target domains but often fails to preserve the semantic structure of the novel classes (Table 3), we purposefully push the novel classes away from the support of the known classes.

We build on the open-set adversarial domain adaptation objective, \mathcal{L}_{adv} , from [41] for domain alignment. However, this \mathcal{L}_{adv} often struggles with fine-grained separation between novel and known classes. To address this limitation, we introduce a margin objective $\mathcal{L}_{\text{margin}}$ that ensures each sample $x \in \mathcal{D}_{\text{syn}}^{e_g}$ is either confidently assigned to a known class or identified as an open-set sample with high probability, improving fine-grained differentiation between known and novel classes. The combined domain alignment loss, \mathcal{L}_{da} , is defined as follows, where λ is a weighting factor:

Methods	Venue	Synthetic-Domains	PACS			Office-Home			DomainNet			Average		
			All	Old	New	All	Old	New	All	Old	New	All	Old	New
VIT [12]	ICLR'21	✗	41.98	50.91	33.16	26.17	29.13	21.62	25.35	26.48	22.41	31.17	35.51	25.73
GCD [49]	CVPR'22	✗	52.28	62.20	38.39	52.71	54.19	50.29	27.41	27.88	26.13	44.13	48.09	38.27
SimGCD [61]	ICCV'23	✗	34.55	38.64	30.51	36.32	49.48	13.55	2.84	2.16	3.75	24.57	30.09	15.94
CMS [10]	CVPR'24	✗	28.95	28.13	36.80	10.02	9.66	10.53	2.33	2.40	2.17	13.77	13.40	16.50
CDAD-Net [40]	CVPR-W '24	✗	69.15	69.40	68.83	53.69	57.07	47.32	24.12	23.99	24.35	48.99	50.15	46.83
SODG-NET [3]	WACV'24	✗	37.43	40.28	28.38	36.53	49.42	14.58	27.77	27.98	26.15	33.91	39.23	23.04
GCD with Synthetic [49]	CVPR'22	✓	65.33	67.10	64.42	50.50	51.48	48.96	24.71	24.80	21.94	46.85	47.78	45.11
SimGCD with Synthetic [61]	ICCV'23	✓	39.76	43.76	35.97	35.57	48.58	12.89	2.71	1.99	4.14	26.01	31.44	17.67
CMS with Synthetic [10]	CVPR'24	✓	28.01	26.71	29.04	12.09	12.66	11.13	3.22	3.28	3.03	14.44	14.22	14.40
CDAD-Net with Synthetic [40]	CVPR-W '24	✓	60.76	61.67	59.49	53.49	56.90	47.76	23.85	23.88	24.26	46.03	47.47	43.84
DAML [43]	CVPR '21	✓	40.26	42.90	29.28	36.20	49.53	14.06	27.10	28.25	26.16	34.52	40.23	23.17
DG ² CD-Net with TIES-Merging [67]		✓	67.04	71.25	64.02	53.52	57.12	48.03	28.72	30.35	24.39	49.76	52.91	45.48
DG ² CD-Net with baseline TA[21]		✓	71.02	73.44	68.01	52.63	52.41	52.69	28.12	29.58	24.42	50.59	51.81	48.37
DG ² CD-Net (Ours)		✓	73.30	75.28	72.56	53.86	53.37	54.33	29.01	30.38	25.46	52.06	53.01	50.78
DG ² CD-Net * (Ours) [LoRA [20]]		✓	75.21	76.31	75.28	54.32	53.03	56.03	27.17	27.61	25.99	52.23	52.32	52.43
CDAD-Net (DA) [40] [Upper bound]		✗	83.25	87.58	77.35	67.55	72.42	63.44	70.28	76.46	65.19	73.69	78.82	68.66

Table 1. **Performance comparison** across three datasets, averaged over all domain combinations. We baseline with two versions of existing GCD and CD-GCD methods: one trained solely on \mathcal{D}_S and another on $\mathcal{D}_S \cup \mathcal{D}_{\text{syn}}$, both without target-specific loss functions to simulate the DG scenario. Subsequently, multiple model aggregation strategies are evaluated for DG²CD-Net. As an upper bound, we include the full CDAD-Net, designed for the transductive DA setting. Red cells highlight the top results; yellow is the second-best.

$$\mathcal{L}_{\text{da}} = \mathcal{L}_{\text{adv}} + \lambda \mathcal{L}_{\text{margin}} \quad (6)$$

\mathcal{L}_{adv} uses a $|\mathcal{Y}_s^{e_g}|+1$ -class episode-specific classifier $\mathcal{F}_c^{e_g}$ on top of \mathcal{F}^{g-1} , with an additional class representing the open-set (details in Section G). For $\mathcal{L}_{\text{margin}}$, let $p(x) = [p_1(x), \dots, p_{|\mathcal{Y}_s^{e_g}|}(x)]$ denote predicted probabilities for known classes by $\mathcal{F}_c^{e_g}$ for x . We enforce a margin m ($0 < m < 1$) between the highest known-class probability, $\max(p(x))$, and the open-set probability, $(1 - \sum_q p_q(x))$:

$$\mathcal{L}_{\text{margin}} = \mathbb{E}_{x \in \mathcal{P}(\mathcal{D}_{\text{syn}}^{e_g})} \max \left(0, m - \left| \max(p(x)) - \left(1 - \sum_{q=1}^{|\mathcal{Y}_s^{e_g}|} p_q(x) \right) \right| \right) \quad (7)$$

$\mathcal{L}_{\text{margin}}$ enhances alignment between domains while preserving known-novel class distinctions. The model is fine-tuned while **minimizing** $\mathcal{L}_{\text{con}}^s + \mathcal{L}_{\text{con}}^u + \mathcal{L}_{\text{da}}$ given $[\mathcal{F}^{g-1}, \mathcal{F}_c^{e_g}]$ to obtain $\theta_{\text{local}}^{e_g}$. Additionally, we examine resource-efficient fine-tuning using LoRA [20], which reduces memory and computational requirements by adding trainable low-rank matrices to the existing weights. During inference, we employ the latest $\theta_{\text{global}}^{n_g}$ and apply clustering on samples from $\mathcal{D}_{\mathcal{T}}$. For the **pseudo-code**, see Section H

4. Experimental Evaluations

Datasets: We evaluate the cross-domain capabilities of DG²CD-Net on three benchmark datasets commonly used for DG [79] and CD-GCD [49]: (i) PACS [25] (4 domains, 9,991 samples, 7 classes), (ii) Office-Home [50] (4 domains, 15,588 samples, 65 classes), and (iii) DomainNet [36] (6 domains, 586,575 samples, 345 classes). For PACS and Office-Home, each domain alternates as the source, with the remaining domains as targets. Source-target configurations for DomainNet are detailed in the Section D.

The known-to-novel class ratios are 4:3 for PACS, 40:25 for Office-Home, and 250:95 for DomainNet, following [40].

Implementation details: All methods employ the global model \mathcal{F} with a ViT-B/16 backbone [12], initialized with DINO pre-trained weights [6], using the [CLS] token as the feature representation. Following the baseline in [49], only the final block of the vision transformer is fine-tuned in each local task model, with an initial learning rate of 0.1, using an SGD optimizer with cosine annealing decay. We set the batch size to 128 and $\lambda = 0.20$. We consider $n_g = 10$, $n_e = 6$, and each episode-specific model is fine-tuned for eight epochs. To balance separation between known and novel classes, we set the margin $m = 0.7$ in the margin loss $\mathcal{L}_{\text{margin}}$ (see Table 15 for ablation). As an alternative, we also fine-tune local models using a LoRA variant with $\alpha = 16$ and $\text{rank} = 8$ for enhanced performance.

We adopt the K -estimation method described in [49] for determining the optimal number of clusters K . This involves applying Brent’s algorithm to features extracted from benchmark datasets, with K constrained between $|\mathcal{Y}_s|$ and 1000 classes, ensuring efficiency and practicality.

Evaluation: Our primary evaluation metric, All, assesses clustering performance across the full $\mathcal{D}_{\mathcal{T}}$. We also report results for two subsets: Old, containing instances from known classes ($\mathcal{Y}_t^{\text{old}}$), and New, containing instances from novel classes ($\mathcal{Y}_t^{\text{new}}$). Subset-specific accuracies are computed using the Hungarian optimal assignment algorithm [24] to determine the label permutation that maximizes accuracy. All reported results reflect the average performance across all domain combinations based on three runs.

4.1. Comparison to the literature

We perform a comparative analysis of DG²CD-Net against DG and category discovery benchmarks. For single-source open-set DG, we include SODG-Net [3]. From the GCD literature, we evaluate ViT [12], GCD [49], SimGCD [61],

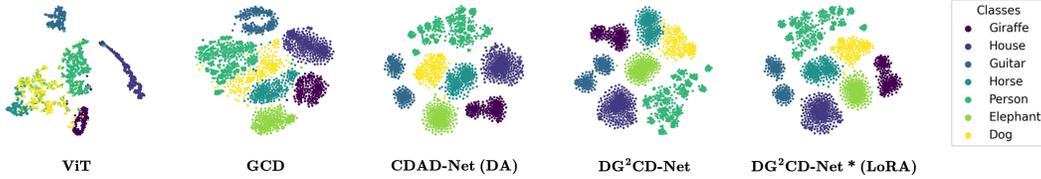


Figure 6. **t-SNE [47] visualizations** of the target domain (“Photo”) clusters, as produced by pre-trained ViT, GCD [49], CDAD-Net(DA) [40], DG²CD-Net and DG²CD-Net* (LoRA) for the PACS dataset, with “Sketch” as the source domain. Both the variants of DG²CD-Net are able to produce a clean and compact embedding space.

CMS [10], and CDAD-Net [40] as the representative of CD-GCD. All models are trained solely on source data, excluding target-specific loss functions to simulate DG-GCD. For GCD, SimGCD, CMS, and CDAD-Net, we also consider variants trained on the combined dataset $\mathcal{D}_S \cup \mathcal{D}_{syn}$. Additionally, we include the multi-source open-set DG method DAML [43]. CDAD-Net (DA) serves as an upper bound for domain adaptation, trained on both source and target domains. Results of CROW (DA) [60], which can also serve as an upper bound, are mentioned in the Table 12.

Results on PACS, Office-Home, and DomainNet are presented in Table 1 (full results in Section J.2). The baseline ViT achieves an average A_{11} score of 31.17%. Among GCD and CD-GCD baselines, CDAD-Net (DG) achieves the highest performance (average A_{11} : 48.99%) when trained on \mathcal{D}_S , demonstrating robustness in domain generalization without synthetic data. In contrast, vanilla GCD [49] performs best (A_{11} : 46.85%) when trained on $\mathcal{D}_S \cup \mathcal{D}_{syn}$, though the mixed distributions introduce gradient conflicts, affecting stability. CDAD-Net (DA) achieves an upper bound A_{11} score of 73.69% in a transductive setting. SODG-Net and DAML trail behind DG²CD-Net by 18.15% and 17.54% in A_{11} , particularly struggling on the New metric due to ineffective clustering of novel classes.

DG²CD-Net achieves 49.76% accuracy using only top-ranked model weights [67], which improves to 50.59% with full-weight-space TA [21], suggesting that the full model space reduces performance ambiguity from gradient conflicts. Our episodic updates with validation-based weighting further boost performance to 52.06%. With LoRA [20] replacing full fine-tuning, DG²CD-Net* reaches 52.23%, indicating that domain-general knowledge resides in a compact manifold across datasets. t-SNE visualizations (Fig. 6) show DG²CD-Net forming a compact embedding space.

4.2. Estimating the number of clusters

Table 2 highlights DG²CD-Net’s accuracy in estimating the number of clusters for PACS, Office-Home, and DomainNet datasets. DG²CD-Net aligns perfectly with the ground truth for PACS at 7 clusters and closely matches it for Office-Home and DomainNet with 67 and 355 clusters, respectively, outperforming CDAD-Net (DG) and achieving parity with CDAD-Net (DA). Such results illustrate our episodic training approach’s strength in adapting to dynamic seman-

tic shifts in \mathcal{D}_T .

Dataset	PACS	Office-Home	DomainNet
Ground Truth	7	65	345
CDAD-Net (DG)	12	60	362
CDAD-Net (DA)	7	66	349
Ours	7	67	355

Table 2. **Estimation of cluster numbers** in target domains showcases DG²CD-Net’s effectiveness in achieving near-optimal clustering, beating other DG counterparts, and closely matching CDAD-Net (DA) performance.

4.3. Main ablation study

Besides the following experiments, additional ablation analyses concerning *different splits of known-novel classes, different backbones, sensitivity to m in Eq. 7, model ablation on more datasets* etc. are mentioned in Section L.

(a) Impact of the number of episodes / synthetic domains in each global update: We evaluate our model’s performance across varying numbers of training episodes, from 0 (using only the source domain) to 12, where each episode introduces a distinct synthetic domain. As shown in Fig. 7, training with up to 6 episodes significantly improves model accuracy, highlighting the impact of episodic learning for domain generalization. Beyond six episodes, however, gains diminish, likely due to increased sign conflicts—where opposing updates from different episodes lead to parameter divergence [67]. This rise in sign conflicts, as visualized in Fig. 3, correlates with the diminishing returns in accuracy, indicating a critical trade-off between generalization and stability. Optimizing the number of episodes is, therefore, vital to balance these effects.

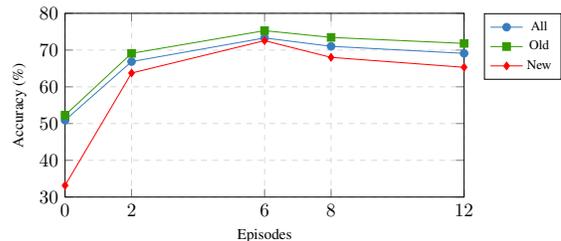


Figure 7. Charting the relationship between **number of episodes and model accuracy** on PACS, this graph shows accuracy peaks at 6 episodes, which is majorly consistent across all the datasets.

(b) Impact of model components: We conducted experiments with several variants of our algorithm to underscore

the impact of key components on domain generalization: (i) **Manual augmentations for \mathcal{D}_{syn}** : synthetic data generated through Gaussian blur, color jittering, and random rotations; (ii) **Without synthetic domain**: training solely on the source domain without target-specific loss (no CD-GCD locally); (iii) **Without multi-global updates**: using a single global update as in conventional Task Arithmetic (TA) [21]; (iv) **Static known/novel splits**: fixing class splits across episodes to limit generalization over semantic shifts; (v) **Conventional normalization of w_g^e** : using min-max instead of softmax normalization in Eq. 2; (vi) **Episode-specific validation set $\mathcal{D}_{\text{valid}}$** : per-episode validation data ($\mathcal{D}_S^{e_g} \cup \mathcal{D}_{\text{syn}}^{e_g}$) instead of a separate $\mathcal{D}_{\text{valid}}$; (vii) **Replacing losses with CDAD-Net** [40]: substituting CDAD-Net’s closed-set domain alignment loss for CD-GCD with ours; and (viii) **DG²CD-Net with Fisher-merging** [30]: incorporating Fisher-merging for model weighting.

As shown in Table 3, removing synthetic domains reduces the All metric by 22.39%, while replacing them with manual augmentations decreases it by 9.38%. A single global update instead of iterative updates lowers the All metric by 6.60%, and static class splits lead to a 4.70% decrease in performance. Min-max normalization reduces All by 11.32%, and episode-specific validation decreases it by 4.04%, indicating the need for out-of-distribution validation data. Replacing our alignment and feature learning losses with CDAD-Net yields a 3.48% drop in All. Lastly, Fisher-merging inadequately captures domain variations, increasing gradient conflicts (Fig. 3) and causing a 1.63% decline in All. These findings underscore the importance of synthetic domains, iterative updates, and our alignment techniques for robust domain generalization.

Model Variant	PACS		
	All	Old	New
(i) With manual augmentations based \mathcal{D}_{syn}	63.92	62.09	67.06
(ii) Without synthetic domain	50.91	52.28	33.16
(iii) Without multi-global updates	66.70	67.52	66.39
(iv) Static known/novel class split across episodes	68.60	69.70	66.13
(v) Conventional normalization for w_g^e in Eq. 2	61.98	69.77	54.43
(vi) Episode specific $\mathcal{D}_{\text{valid}}$	69.26	73.34	65.15
(vii) Replacing our CD-GCD losses with those of [40]	69.82	74.00	66.23
(viii) DG ² CD-Net with Fisher-merging [30]	71.67	74.42	70.03
Full DG ² CD-Net	73.30	75.28	72.56

Table 3. Performance metrics demonstrating the **influence of key model components** of DG²CD-Net for PACS.

(c) **Impact of loss components**: Table 4 demonstrates the substantial improvements in classification accuracy on the PACS dataset achieved by systematically integrating loss components. Implementing $\mathcal{L}_{\text{con}}^s$ in **C-1** raised accuracy to 60.63%, showcasing its efficacy with labeled data. With $\mathcal{L}_{\text{con}}^u$ in **C-2** accuracy boosted to 67.84% through the use of unsupervised data. Incorporating \mathcal{L}_{adv} in **C-6** elevated the accuracy to 71.45%. Finally, the complete integration in **C-7**, incorporating $\mathcal{L}_{\text{margin}}$, reached 73.30%. These results underscore the synergistic effect of the combined loss functions, significantly enhancing the model’s generaliza-

Configurations	$\mathcal{L}_{\text{con}}^s$	$\mathcal{L}_{\text{con}}^u$	\mathcal{L}_{adv}	$\mathcal{L}_{\text{margin}}$	PACS		
					All	Old	New
C-0	X	X	X	X	41.98	50.91	33.16
C-1	✓	X	X	X	60.63	65.15	54.63
C-2	X	✓	X	X	67.84	68.92	66.17
C-3	✓	✓	X	X	68.70	69.52	67.61
C-4	✓	X	✓	X	64.16	66.10	60.81
C-5	X	✓	✓	X	69.22	69.20	69.34
C-6	✓	✓	✓	X	71.45	72.92	70.33
C-7	✓	✓	✓	✓	73.30	75.28	72.56

Table 4. Impact of **loss components** of DG²CD-Net on PACS.

tion capability.

4.4. Analyzing training convergence

Fig. 8 illustrates the progressive weight difference norms for θ_{global} ($|\theta_{\text{global}}^g - \theta_{\text{global}}^{g-1}|_1$) across ten global updates in the attention projection weights and MLP layers of our ViT-B/16 block, trained on the “Photo” domain of the PACS dataset. The results show a trend toward stabilized weight differences, indicating a smooth model convergence.

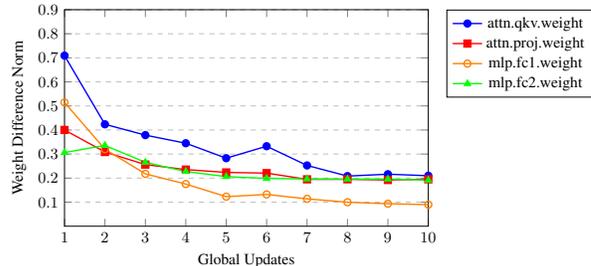


Figure 8. Plot of the **average weight difference norm across global updates for different layers**, indicating model convergence and stabilization.

5. Conclusions & Future Directions

This paper introduces DG-GCD, a novel GCD setting where the unlabeled target domain, distinct from the labeled source domain, is accessible only at test time. To address clustering under domain shifts without target-specific data, we propose DG²CD-Net, which integrates episodic training on source and synthetic pseudo-target domains with adaptive task-vector aggregation for a robust global model. Unlike traditional Task Arithmetic, our method dynamically weights task vectors based on the generalization performance of episode-specific models on held-out validation distributions. Additionally, an alignment objective with a novel margin loss ensures effective knowledge transfer, enabling confident differentiation of known and novel classes within episodes. Comprehensive benchmarks and ablation studies across multiple datasets validate the effectiveness of DG²CD-Net. We sincerely hope DG-GCD inspires further research in challenging areas like spotting new identities in person re-identification and novel drug discovery from chemical structures, where prior knowledge is scarce.

Acknowledgment. We acknowledge Adobe Research, CMInDS IITB for computational resources and our colleague Sona Elza Simon for insightful discussions.

Supplementary Material

A. What's Inside This Supplementary Material?

This supplementary material provides additional insights and extended analysis to support the main paper. Below, we outline the key sections :

Extended Literature on Class Discovery: ([Section B](#)) A detailed review of prior work in Novel Category Discovery (NCD) and Generalized Category Discovery (GCD), emphasizing advancements relevant to domain shifts.

DG-GCD in Practice: Applications Across Domains: ([Section C](#)) Real-world applications of DG-GCD, including driverless cars, healthcare, and retail.

Dataset Details: ([Section D](#)) Comprehensive information on datasets, their configurations, and class distributions.

Technical Details about Synthetic Data Generation: ([Section E](#)) Description of the synthetic domain generation pipeline and parameter configurations.

Synthetic Domain Utilization in DG-GCD: ([Section F](#)) Integration of synthetic domains into training and validation, with visualizations and statistical insights.

Handling Open-Set Recognition: ([Section G](#)) Methodologies for open-set recognition, including episodic classifiers and adversarial loss.

Pseudocode for Episodic Training Strategy: ([Section H](#)) The pseudocode for our episodic training strategy and its implementation.

Meta-Knowledge Learned in Episodic Training: ([Section I](#)) An exploration of how meta-knowledge is acquired and refined through episodic training, enabling robust domain generalization and task adaptability.

Comprehensive Comparative Analyses Across Datasets: ([Section J.2](#)) Comparative analysis of DG²CD-Net on PACS, Office-Home, and DomainNet, evaluating its robustness and adaptability against benchmarks.

Performance Comparison with Domain Adaptation Methods: ([Section K](#)) A detailed comparison of DG²CD-Net with baseline and upper-bound domain adaptation methods.

Additional Ablation Studies: ([Section L](#)) An analysis of the contributions of individual components in DG²CD-Net across multiple datasets.

Effect of Backbone Initialization: ([Section M](#)) A comparison of results using different backbones.

Effect of LoRA Fine-Tuning: ([Section N](#)) A comparison of results using DG²CD-Net with different LoRA adapters.

Limitations and Future Work: ([Section O](#)) A discussion on the current limitations and potential future directions for DG²CD-Net.

Each section is carefully crafted to provide deeper insights, reproducibility details, and additional context for the results presented in the main paper. We hope this supplementary material enhances the reader's understanding and offers a foundation for future exploration of domain generalization and category discovery.

B. Extended Literature on Class discovery

Category discovery has evolved significantly from Novel Category Discovery (NCD) [16] to GCD [49]. Traditionally, NCD methods [16, 17, 73, 75, 76] have utilized dual-model architectures where separate models are trained on labeled and unlabeled data to facilitate task transfer or employed parametric classifiers on top of generic feature extractors to categorize new classes. Recent advancements in GCD focus on leveraging labeled data to generate pseudo-labels for unlabeled images. DCCL [38] uses InfoMap clustering, while PromptCAL [72] identifies pseudo-positive samples through semi-supervised affinity propagation techniques. PIM [9] improves this by optimizing bi-level mutual information, and SimGCD [61] utilizes knowledge distillation with a parametric classifier to enhance pseudo-label reliability. A key innovation is the contrastive mean-shift clustering from [10], which creates a highly discriminative embedding space for fine-grained class distinctions. Additionally, Gaussian mixture models [74] have been explored for clustering in GCD.

Methods	$\mathcal{D}_{\mathcal{T}}$ Available for Training	Evaluation on Known/Novel classes	$\mathcal{P}(\mathcal{S}) = \mathcal{P}(\mathcal{T})$
NCD [15, 19]	✓	Novel	✓
GCD [49]	✓	Both	✓
CD-GCD [40]	✓	Both	✗
DG-GCD	✗	Both	✗

Table 5. Comparison of methods based on target-domain data ($\mathcal{D}_{\mathcal{T}}$) availability during training, evaluation on known and/or novel classes, and distribution divergence between the source (\mathcal{S}) and target (\mathcal{T}) domains in class discovery contexts: novel (NCD), generalized (GCD), and cross-domain (CD-GCD). Our proposed DG-GCD setting is different and more challenging than the rest.

C. DG-GCD in Practice: Applications Across Domains

The challenge of *domain generalization for generalized category discovery (DG-GCD)* has numerous real-world applications across various industries. In *autonomous vehicles*, models must adapt to changing environmental conditions (e.g., weather, lighting) and detect novel road objects without access to original training data. In *healthcare*, medical image analysis must generalize across different diagnostic tools while discovering new pathologies, all while respecting patient privacy regulations. *Wildlife monitoring* requires models to generalize across ecosystems while identifying new species, while *surveillance systems* must detect novel threats and adapt to different environments such as airports and public spaces. In *retail and e-commerce*, recommendation systems need to adapt across product categories and discover new items, while *robotics* requires models to generalize in dynamic environments and detect new objects. Lastly, in *precision agriculture*, systems must generalize across different farms and detect novel crop diseases without requiring access to original datasets. These applications highlight the importance of models that can generalize across domains and discover new categories in privacy-preserving, real-world scenarios.

D. Dataset details

Datasets: Our experiments are conducted on three benchmark datasets (i) PACS [25] (ii) Office-Home [50], and (iii) DomainNet [36].

For PACS and Office-Home, each domain was used as the source, with all others as target domains. For DomainNet, a subset of source-target configurations was selected, as summarized in Table 6, ensuring the model was tested on diverse domain pairs.

Source	Targets
Sketch	Clipart, Painting, Infograph, Quickdraw, Real World
Painting	Clipart, Sketch, Infograph, Quickdraw, Real World
Clipart	Painting, Sketch, Infograph, Quickdraw, Real World

Table 6. Source-target configurations for DomainNet

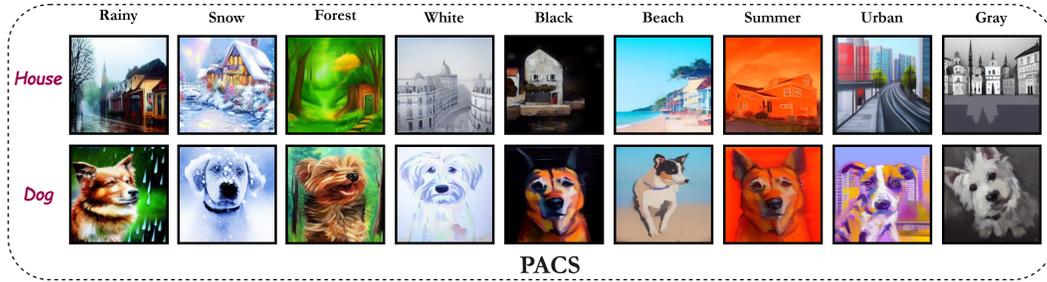


Figure 9. Additional synthetic image samples generated for the PACS dataset, showcasing diverse styles and domains across different categories (House and Dog).

E. Technical details about Synthetic Data Generation

We employed the InstructPix2Pix pipeline from Hugging Face’s Diffusers library, utilizing the `timbrooks/instruct-pix2pix` model, for image transformation tasks. To achieve a balance between processing time and output quality, we configured the key parameter `num_inference_steps` to 10. The `image_guidance_scale` parameter was set to 1.0, ensuring that the model retained the essential structure of the input image while applying the specified transformations. Furthermore, the `guidance_scale` parameter was adjusted to 7.5, promoting a strong alignment with the transformation prompt. These configurations allow for straightforward replication of our process while maintaining high-quality output.

F. Synthetic Domain Utilization in DG-GCD

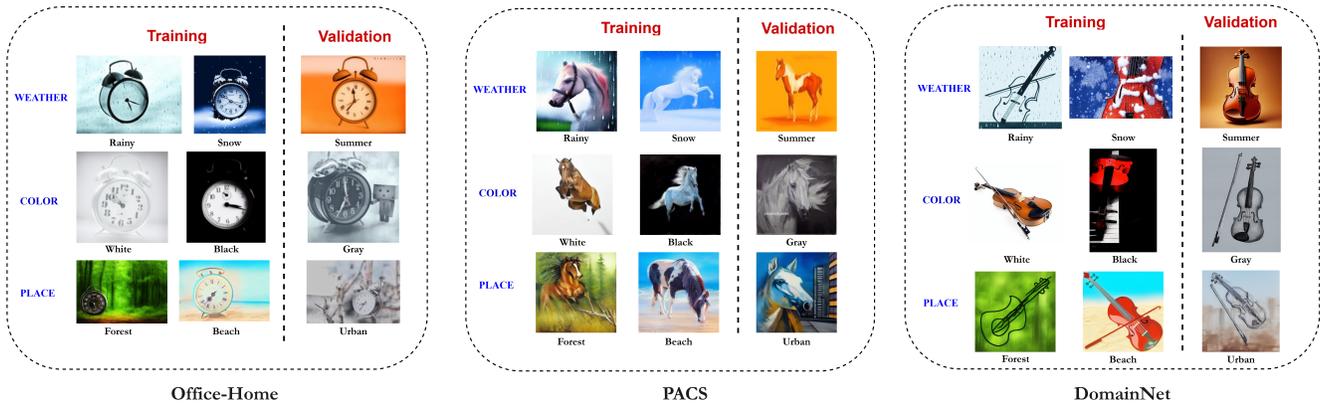


Figure 10. Categorization of synthetic domains utilized in the training and validation phases. Training domains are designed to simulate diverse conditions such as weather, color, and place variations. Validation domains challenge the model’s adaptability to new, complex scenarios.

In our study, we generate nine synthetic domains for each of the PACS, Office-Home, and DomainNet datasets to enhance the adaptability of models trained under the Domain Generalization for Generalized Category Discovery (DG-GCD) framework. Table 7 outlines the distribution of these domains, with six utilized in training and three in validation, ensuring comprehensive exposure to varied environmental conditions and rigorous testing of generalization capabilities. This structured approach improves the models’ robustness against unseen real-world scenarios.

Figure 10 categorizes the synthetic domains used in our study’s training and validation phases. The training domains include variations such as “Rainy” and “Snow” weather, “White” and “Black” colors, and “Forest” and “Beach” settings, broadening the model’s exposure to diverse scenarios. The validation domains introduce “Summer” weather, “Gray” color, and “Urban” settings to test the model’s ability to generalize across new and complex environments. This strategic use of synthetic domains demonstrates our approach to enhancing robustness and adaptability in models, crucial for effective domain generalization in real-world applications.

Dataset	Total Domains Generated	Used in Training	Used in Validation
PACS	9	6	3
Office-Home	9	6	3
DomainNet	9	6	3

Table 7. Synthetic Domain Generation and Utilization for DG-GCD

Table 8 presents the Average Fréchet Inception Distance (FID) scores for synthetic domains compared to the original domains within the Office-Home dataset. The scores are calculated to evaluate the visual similarity between generated images in environments like “Rainy,” “Black,” “Urban,” “Beach,” “White,” and “Snow,” and the original dataset categories: Art, Clipart, Product, and Real World. Lower FID scores indicate closer visual resemblance to the original domain, suggesting better synthetic image quality and domain adaptation. The data reveal variations in FID scores, with “White” and “Black” environments achieving the lowest scores, indicating higher similarity and potentially more effective domain generalization. This analysis provides insights into which synthetic modifications most accurately reflect the characteristics of their respective real-world counterparts, crucial for training robust models capable of generalizing across diverse visual contexts.

Domain	FID Scores			
	Art	Clipart	Product	Real World
Rainy	140.11	158.81	167.27	154.42
Black	58.69	74.35	86.65	72.28
Urban	97.77	131.80	138.45	119.78
Beach	109.13	133.33	136.54	121.10
White	49.80	71.07	77.34	61.73
Snow	100.99	125.03	132.42	117.50

Table 8. Average Fréchet Inception Distance (FID) comparison between pairs of the generated and original domains over all the classes of the Office-Home dataset.

G. Handling Open-Set Recognition with Episode-Specific Classifiers

In our framework, we employ a $|\mathcal{Y}_s^{e_g}|+1$ -class episode-specific classifier to address open-set recognition in episodic training. Here, $|\mathcal{Y}_s^{e_g}|$ denotes the number of known classes in an episode, while the additional class models the open-set category, capturing instances outside the known classes. This design dynamically adapts to episodic data, ensuring robust differentiation between known and unknown classes.

G.1. Adversarial Loss in Open-Set Domain Adaptation

To refine the decision boundary, we use an adversarial loss, \mathcal{L}_{adv} , which separates open-set instances by pushing them further from the known classes in feature space. This adversarial refinement enables the confident classification of unknown samples while maintaining performance on known categories. By tailoring decision boundaries to the episodic data, the classifier improves adaptive learning, generalization to new domains, and detection of novel classes in unseen distributions.

The adversarial loss used in our method is inspired by [41]. It facilitates the separation of known and unknown samples in the target domain by training the classifier $\mathcal{F}_c^{e_g}$ and the generator \mathcal{F}^{g-1} adversarially. The adversarial loss \mathcal{L}_{adv} is defined as:

$$\mathcal{L}_{adv}(x^t) = -\alpha \log(p(y = |\mathcal{Y}_s^{e_g}|+1|x^t)) - (1 - \alpha) \log(1 - p(y = |\mathcal{Y}_s^{e_g}|+1|x^t)), \quad (8)$$

Where:

- x^t represents a target sample from \mathcal{D}_{syn} ,
- $p(y = |\mathcal{Y}_s^{e_g}|+1|x^t)$ is the predicted probability that x^t belongs to the unknown class,
- α is a hyperparameter (set to 0.5 in our experiments) that determines the decision boundary for the unknown class.

G.1.1. Training Objectives

Adversarial loss (\mathcal{L}_{adv}) is optimized along with Source Classification loss, which is a standard cross-entropy loss that ensures accurate classification of known source samples:

$$\mathcal{L}_s(x^s, y^s) = -\log(p(y = y^s|x^s)), \quad (9)$$

where (x^s, y^s) represents a source sample and its label from \mathcal{D}_S respectively.

The classifier \mathcal{F}_c^{eg} and generator \mathcal{F}^{g-1} is trained in the following manner:

- For the classifier \mathcal{F}_c^{eg} , the objective is to minimize the total loss:

$$\min_{\mathcal{F}_c^{eg}} \mathcal{L}_s(x^s, y^s) + \mathcal{L}_{adv}(x^t). \quad (10)$$

- For the generator \mathcal{F}^{g-1} , the objective is to deceive the classifier by maximizing the adversarial loss:

$$\min_{\mathcal{F}^{g-1}} \mathcal{L}_s(x^s, y^s) - \mathcal{L}_{adv}(x^t). \quad (11)$$

G.1.2. Implementation Details

To efficiently compute the adversarial loss, we use a Gradient Reversal Layer (GRL), which flips the gradient sign during backpropagation. This allows simultaneous updates to \mathcal{F}_c^{eg} and \mathcal{F}^{g-1} , facilitating stable training of adversarial objectives.

H. Pseudocode of the proposed Episodic Training Strategy

In this section, we present the pseudocode for the proposed episodic training strategy, as detailed in the main paper. This algorithm is used to iteratively update the global model parameters θ_{global} across multiple episodes and global updates. The process involves training task models on synthetic domains and updating the global model based on task vector computations and validation results.

I. Meta-Knowledge learnt in Episodic Training

In our episodic training framework, meta-knowledge encompasses the cumulative insights gained from dynamic adaptation to varying domain conditions. This knowledge is perpetually refined via systematic application and iterative adjustment of task vectors, informed by feedback from domain-specific fine-tuning and rigorous validation processes. Unlike conventional meta-learning, which primarily targets rapid task adaptability, our framework emphasizes robust domain generalization. This approach enhances the model’s proficiency in effectively preemptively addressing and adapting to evolving data distributions.

Meta-knowledge is acquired through:

- [-] **Cross-Domain Exposure:** By engaging with overlapping features across multiple domains, the model develops a nuanced capability to generalize across diverse training environments. This cross-domain learning is fundamental in enabling the model to abstract and apply domain-invariant patterns to new, unseen scenarios.
- [-] **Dynamic Vector Adjustments:** Task vectors are continually updated in response to real-time performance metrics. This dynamic refinement process allows the model to adjust its generalization strategies on the fly, enhancing its responsiveness to changes in domain characteristics.
- [-] **Validation-Driven Learning:** The integration of internal validation mechanisms ensures continuous performance feedback. This feedback is instrumental in fine-tuning the model’s strategic adjustments, ensuring optimized responses to future domain shifts and data interactions.

Our method enhances the generalizability of pre-trained foundation models by adaptively combining their fine-tuned versions across multiple automatically synthesized domains, eliminating the need for manual annotations. This adaptive strategy leverages the generalization performance of each fine-tuned model, minimizing the impact of poorly generalizable models. Besides, our approach ensures both discriminativeness and domain independence for the DG-GCD task. Consequently, this produces an embedding space predominantly guided by class semantics and suppressing stylistic artifacts, making it highly effective for generalization and clustering.

J. More details on comparisons to literature

J.1. DG-GCD specific adaptations for baselines

We evaluate several state-of-the-art methods for generalized category discovery (GCD) and domain generalization (DG) by adapting them to the DG-GCD setting, as detailed in Table 2 of the main text. In this setting, target domain access

Algorithm 1 Proposed Episodic Training Strategy for Updating θ_{global}

Require: Pre-trained global model parameters θ_{global}^0 , number of global updates n_g , number of episodes per global update n_e , source domain $\mathcal{D}_S^{e_g}$, synthetic domain $\mathcal{D}_{\text{syn}}^{e_g}$, validation domain $\mathcal{D}_{\text{valid}}$ in the e^{th} episode.

Ensure: Final global model $\theta_{\text{global}}^{n_g}$.

- 1: **for** $g = 1$ to n_g **do** ▷ Global updates
- 2: Randomly shuffle synthetic domains $\mathcal{D}_{\text{syn}}^{e_g}$.
- 3: **for** each episode $e = 1$ to n_e **do** ▷ Episode training
- 4: Initialize task model parameters $\theta_{\text{local}}^{e_g} \leftarrow \theta_{\text{global}}^{g-1}$.
- 5: Train task model $\theta_{\text{local}}^{e_g}$ on $(\mathcal{D}_S^{e_g}, \mathcal{D}_{\text{syn}}^{e_g})$ for the CD-GCD task.
- 6: Compute task vector δ_g^e (Equ. 1 in the main text) :

$$\delta_g^e = \theta_{\text{global}}^{g-1} - \theta_{\text{local}}^{e_g}$$

- 7: Validate task model on $\mathcal{D}_{\text{valid}}$ and compute accuracies: All, Old, New.
- 8: **end for**
- 9: Compute weights w_g^e using softmax on All accuracies (Equ. 2 in main text) :

$$w_g^e = \frac{\exp(\text{All}_g^e)}{\sum_{e'=1}^{n_e} \exp(\text{All}_g^{e'})}$$

- 10: Update global model θ_{global}^g (Equ. 3 in main text) :

$$\theta_{\text{global}}^g = \theta_{\text{global}}^{g-1} - \sum_{e=1}^{n_e} w_g^e \delta_g^e$$

11: **end for**

- 12: Save the final global model: $\theta_{\text{global}}^{n_g}$.
-

is completely removed during training, and for certain methods, synthetic domain data is incorporated to simulate domain shifts.

For ViT-B/16, pre-trained with DINO, we fine-tuned only the last block using source domain data, following standard GCD practices, and evaluated the model on target domains without any target domain data during training. Similarly, GCD was adapted by fine-tuning the last block of the backbone on the source domain alone, and we introduced a synthetic domain variant to account for domain shifts.

For CMS (Contrastive Mean Shift) and SimGCD, we followed a similar procedure. We fine-tuned the last block using only the source domain and, in addition, created synthetic variants by incorporating synthetic domain data to assess the methods' ability to handle domain shifts effectively.

CDAD-Net, which is designed for cross-domain adaptation, was also adapted for the DG-GCD setting. We ensured that it trained solely on the source domain, without access to target domain data, and created a synthetic variant to evaluate its performance on unseen domains.

As seen in Table 2 of the main text, the methods incorporating synthetic domain data, such as GCD with synthetic data, generally performed better in handling domain shifts, especially for novel class discovery. Our proposed models, leveraging task merging techniques such as TIES-Arithmetic and LoRA, outperform baseline methods on both benchmark datasets, achieving superior results, particularly on novel classes. The inclusion of synthetic domains proves beneficial, as evidenced by the marked performance improvement across all datasets, with our model consistently achieving the highest or second-highest results.

J.2. Comprehensive comparative analyses of DG²CD-Net across multiple datasets

This section presents a concise comparative analysis of DG²CD-Net on PACS, Office-Home, and DomainNet in Table-9, 10 and 11 respectively. Each dataset challenges DG²CD-Net with unique domain shifts, showcasing its adaptability and robustness. This evaluation aims to validate DG²CD-Net's performance against established benchmarks, highlighting it's

strengths and identifying opportunities for advancement in domain generalization.

PACS																		
Methods	Art Painting → Sketch			Art Painting → Cartoon			Art Painting → Photo			Photo → Art Painting			Photo → Cartoon			Photo → Sketch		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
ViT [12]	37.44	50.73	19.5	47.4	61.3	35.25	76.05	87.13	64.64	53.17	77.31	31.67	47.01	55.54	39.57	31.87	37.57	24.16
GCD [49]	32.02	41.53	19.12	46.78	60.35	28.57	79.16	99.45	48.73	74.73	80.26	67.31	57.53	60.46	53.6	46.23	48.56	43.08
SimGCD [61]	29.35	17.3	62.12	23.08	28.26	16.32	51.98	74.44	33.26	46.29	48.96	43.17	34.26	44.91	20.35	24.84	31.88	5.68
CDAD-Net [40]	46.02	45.95	46.21	51.71	53.43	49.46	99.04	99.21	98.9	76.61	76.97	76.19	56.78	56.67	56.93	46.65	46.15	48.01
GCD With Synthetic	45.78	36.71	58.01	54.84	73.47	38.57	82.6	66.29	99.39	79	86.84	72.02	53.56	67.93	41.01	44.18	47.78	39.32
CDAD-Net with Synthetic	43.09	42.53	44.6	49.45	59.31	36.58	99.16	99.21	99.12	65.38	62.83	68.36	42.92	41.97	44.15	41.51	43.79	35.32
DG ² CD-Net(TIES-Merging[67])	41.31	40.56	42.31	45.69	57	35.81	96.11	97.87	94.29	62.87	87.72	40.72	48.98	60.02	39.33	44.1	36.33	54.58
DG ² CD-Net(TA[21])	46.4	51.3	42.8	56.21	58.82	52.7	99.2	99.5	98.8	81.47	91.09	68.57	57.76	56.11	59.99	46.88	48.96	44.06
DG ² CD-Net(Ours)	46.79	38.13	58.49	57.96	73.38	44.48	99.34	99.7	98.97	86.67	91.87	82.04	62.97	71.18	55.8	45.72	36.53	58.13
DG ² CD-Net*(Ours)[LoRA[20]]	46.83	37.79	59.03	63.82	71.13	57.43	99.46	99.35	99.57	88.89	93.94	84.4	64.19	72.23	57.15	46.45	37.75	58.19

Methods	Sketch → Art Painting			Sketch → Cartoon			Sketch → Photo			Cartoon → Art Painting			Cartoon → Sketch			Cartoon → Photo		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
ViT [12]	23.93	26.53	21.61	40.61	58.92	24.62	33.29	33.88	32.69	38.09	47.36	29.82	33.57	35.67	30.74	41.38	39.08	43.74
GCD [49]	33.25	39.09	25.43	40.89	48.14	31.17	46.86	59.28	28.22	58.15	78.52	30.86	36	44.83	24.04	75.75	85.88	60.55
SimGCD [61]	21.19	31.91	8.67	23.17	36.77	5.4	34.22	27.46	40.8	38.38	42.07	34.07	34.84	33.94	37.31	53.05	45.85	59.06
CDAD-Net [40]	87.99	84.32	92.28	51.88	51.77	52.02	99.04	99.21	98.9	73.05	76.88	68.57	41.84	42.71	39.49	99.22	99.47	99.01
GCD With Synthetic	82.15	85.13	79.5	44.3	48.22	40.89	99.49	99.76	99.21	63.01	63.73	62.37	35.66	29.95	43.36	99.43	99.47	99.39
CDAD-Net with Synthetic	61.91	69.45	53.12	48.59	53.13	42.67	68.44	63.5	72.56	67.24	65.28	69.52	42.05	39.61	48.67	99.34	99.47	99.23
DG ² CD-Net(TIES-Merging[67])	80.59	80.78	80.42	58.94	75.71	44.28	99.07	98.64	99.51	87.45	90.93	84.35	40.67	31.39	53.2	98.71	97.99	99.45
DG ² CD-Net(TA[21])	73.02	79.37	64.51	55.89	54.84	57.29	99.31	99.5	99.03	90.89	92.75	88.4	46.03	49.67	41.1	99.16	99.35	98.88
DG ² CD-Net(Ours)	88.75	93.52	84.49	56.76	72.14	43.33	99.13	98.7	99.57	90.77	93.37	88.46	49.2	43.18	57.33	95.57	91.62	99.64
DG ² CD-Net*(Ours)[LoRA[20]]	90.87	95.28	86.93	66.25	78.32	55.72	99.22	98.88	99.57	91.02	93.99	88.37	46.33	38.19	57.33	99.22	98.82	99.64

Table 9. Detailed comparison of our proposed DG²CD-Net on DG-GCD with respect to referred literature for PACS Dataset.

Office-Home																		
Methods	Art → Clipart			Art → Product			Art → Real World			Clipart → Art			Clipart → Real World			Clipart → Product		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
ViT [12]	18.88	20.86	15.79	30.34	35.42	21.83	29.52	32.76	24.85	14.96	15.6	14.12	18.59	20.12	16.4	30.39	32.51	26.84
GCD [49]	31.65	32.11	30.93	63.18	64.35	61.22	63.85	66.56	59.96	51.96	52.7	51	62.62	65.29	58.79	60.59	67.13	49.61
SimGCD [61]	24.54	34.35	8.09	41.95	57.92	13.54	46.78	65.54	14.73	31.11	39.56	11.88	25.66	37.66	5.15	28.88	41.38	12.96
CDAD-Net [40]	30.95	33.65	26.43	64.99	68.04	59.32	67.5	70.89	61.72	53.36	56.05	47.23	64.7	69.4	55.25	67.02	68.8	63.7
GCD With Synthetic	29.86	31.04	28.02	57.92	63.12	49.19	59.47	59.59	59.29	53.3	52.84	53.89	61.46	58.27	66.06	63.84	64.04	63.51
CDAD-Net with Synthetic	31.97	35.1	26.71	65.39	68.94	62.51	67.83	70.87	62.64	53.51	56.65	46.37	66.97	69.76	62.2	61.4	65.55	57.4
DG ² CD-Net(TIES-Merging[67])	33.96	37	29.23	59.99	62.93	55.07	66.26	68.42	63.15	52.18	52.3	52.04	58.16	58.62	57.5	65.32	72.33	53.56
DG ² CD-Net(TA[21])	29.52	27.31	33.06	62.42	61.67	63.59	64.46	62.14	67.8	51.24	53.32	47.12	64.23	61.24	68.92	65.28	66.03	64.13
DG ² CD-Net(Ours)	31.51	31.96	30.81	67.46	68.73	65.32	64.45	60.25	70.48	50.76	48.76	53.36	64.77	60.58	70.79	65.34	67.48	61.76
DG ² CD-Net(Ours)[LoRA [20]]	31.56	31.85	31.1	65.22	65.68	64.45	67.81	65.14	71.66	53.4	48.47	59.81	66.13	61.63	72.61	66.16	67.12	64.57

Methods	Product → Art			Product → Real World			Product → Clipart			Real World → Art			Real World → Product			Real World → Clipart		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
ViT [12]	23.2	24.64	21.33	31.21	35.45	25.13	19.27	20.52	17.31	32.22	35.79	27.58	44.67	52.21	32.03	20.8	23.71	16.26
GCD [49]	50.27	48.18	52.99	65.07	63.09	67.91	29.08	29.22	28.87	54.26	54.05	54.55	69.04	72.76	62.79	31.04	34.93	24.97
SimGCD [61]	38.28	50.42	10.66	48.36	67.07	16.41	22.45	32.37	11.34	48.95	66.79	8.36	57.19	69.23	44.15	21.7	31.46	5.33
CDAD-Net [40]	50.1	52.43	44.67	66.47	72.13	56.81	31.36	34.6	25.94	54.68	58.07	46.96	61.39	64.79	55.06	31.78	36.02	24.69
GCD With Synthetic	49.18	46.54	52.61	63.4	59.67	68.77	28.43	27.72	29.55	51.71	61.55	38.91	61.14	65.34	54.1	26.38	28.11	23.68
CDAD-Net with Synthetic	54.12	57.67	46.04	66.97	70.2	61.46	32.34	35.13	28.68	53.72	56.89	46.5	56.47	62.33	45.62	31.19	33.67	27.02
DG ² CD-Net(TIES-Merging[67])	53.28	54.77	51.33	62.74	66.85	56.83	31.82	33.97	28.46	57.11	66.14	45.36	67.04	74.25	54.95	34.41	37.94	28.9
DG ² CD-Net(TA[21])	49.92	52.33	45.17	65.57	67.22	62.99	31.48	30.21	33.51	51.65	55.06	44.92	65.01	63.73	66.99	30.73	28.65	34.08
DG ² CD-Net(Ours)	52.45	50.51	54.98	67.87	69.88	64.97	30.71	30.05	31.75	52.31	49.42	56.07	67.37	71.65	60.19	31.28	31.13	31.51
DG ² CD-Net*(Ours)[LoRA[20]]	52.66	51.75	53.84	65.48	62	70.48	31.52	31.83	31.04	53.42	51.6	55.78	66.33	68.97	61.91	32.26	30.4	35.15

Table 10. Detailed comparison of our proposed DG²CD-Net on DG-GCD with respect to referred literature for Office-Home Dataset

DomainNet															
Methods	Sketch → Real			Sketch → Quickdraw			Sketch → Infograph			Sketch → Painting			Sketch → Clipart		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
VIT [12]	47.17	47.92	44.95	12.13	12.1	12.21	11.99	12.68	10.28	30.95	33.02	25.75	32.64	34.29	28.64
GCD [49]	51.13	51.88	48.92	16.08	15.65	17.2	12.6	12.57	12.68	35.25	35.96	33.46	31.22	30.85	32.1
SimGCD [61]	3.11	3.47	2.32	2.31	2.4	2.1	3.16	2.27	5.24	4.1	2.57	5.62	3.02	2.3	4.07
CDAD-Net [40]	48.21	47.7	49.77	12.27	11.52	14.24	12.07	12.69	11.34	35.47	36.39	32.86	18.63	17.52	20.39
GCD With Synthetic	53	51.71	47.64	13.71	13.79	13.99	12.24	11.99	11.37	35.43	34.12	30.83	22.49	22.2	21.49
CDAD-Net with Synthetic	47.11	46.09	49.4	12.75	13.1	14.05	12.52	13.04	11.92	35.87	36.73	33.35	18.99	17.68	21.07
DG ² CD-Net (TIES-Merging[67])	50.32	52.88	42.8	15.22	15.12	15.49	14.75	16.04	11.53	35.84	38.99	27.93	31.06	33.34	25.53
DG ² CD-Net [TA[21]]	51.84	52.58	49.65	13.67	13.44	14.25	12.72	13.05	11.89	33.96	35.32	30.55	21.94	21.8	22.29
DG ² CD-Net (Ours)	53.67	55.48	48.35	15.9	16	15.63	14.63	15.66	12.06	37.44	39.53	32.19	30.47	32.89	24.58
DG ² CD-Net* (Ours)[LoRA[20]]	53.01	53.75	50.84	13.71	13.38	14.57	13.82	14.23	12.78	36.77	37.9	33.93	24.17	24.46	23.44

Methods	Clipart → Infograph			Clipart → Quickdraw			Clipart → Sketch			Clipart → Real			Clipart → Painting		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
VIT [12]	12.18	12.64	11.03	12.13	12.1	12.21	24.76	26.24	21.27	44.14	45.43	40.34	26.76	28.7	21.91
GCD [49]	14.03	14.64	12.49	14.94	14.67	15.65	25.33	27.68	19.78	53.23	55.48	46.62	34.82	36.82	29.83
SimGCD [61]	2.03	0.4	3.94	0.5	0.3	1	1	0.02	3.842	1.64	1.07	2.42	2.07	2.05	2.13
CDAD-Net [40]	12.79	12.96	12.87	12.06	11.59	12.78	19	19.17	18.76	47.06	44.62	49.2	34.45	36.02	32.85
GCD With Synthetic	11.46	12.03	10.04	12.68	12.57	12.95	18.74	20.54	14.47	50.11	52.26	43.79	32.67	34.91	27.06
CDAD-Net with Synthetic	13	13.37	12.56	12.07	11.76	12.89	17.46	18.03	16.67	48.25	47.51	49.6	33.23	32.79	34.2
DG ² CD-Net (TIES-Merging[67])	15.66	17.02	12.28	14.91	14.73	15.39	27.75	30.64	20.89	54.18	56.73	46.72	36.71	38.14	33.16
DG ² CD-Net [TA[21]]	15.71	16.88	12.78	14.63	14.18	15.81	27.03	29.89	20.26	53.91	55.14	50.29	36.85	39.69	29.75
DG ² CD-Net (Ours)	15.81	17.09	12.63	14.53	14.14	15.58	26.86	29.49	20.64	54.54	56.03	50.17	36.81	38.87	31.67
DG ² CD-Net* (Ours)[LoRA[20]]	14.19	14.12	14.35	13.31	13.23	13.53	22.01	22.9	19.91	53.95	54.99	50.91	37.12	37.89	35.21

Methods	Painting → Infograph			Painting → Quickdraw			Painting → Sketch			Painting → Real			Painting → Clipart		
	All	Old	New	All	Old	New	All	Old	New	All	Old	New	All	Old	New
VIT [12]	12.2	13.1	9.94	12.13	12.1	12.21	23	24.78	18.79	51.53	54.16	43.8	26.57	28.08	22.92
GCD [49]	12.87	12.67	13.37	10.74	10.56	11.21	21.49	22.26	19.68	52.12	51.86	52.86	25.32	24.79	26.6
SimGCD [61]	3.2	2.6	3.8	3.5	2.32	4.65	4.23	3.56	4.86	4.2	3.52	5	4.49	3.6	5.23
CDAD-Net [40]	11.65	12.49	10.66	11.98	11.2	12.44	17.11	17.68	16.32	49.04	48.63	50.27	20.06	19.74	20.57
GCD With Synthetic	10.86	10.56	9.84	11.81	11.8	11.77	17.26	16.25	13.83	49.1	47.3	42.04	19.3	19.45	18.04
CDAD-Net with Synthetic	11.53	12.32	10.59	11.86	10.71	12.32	17.29	18.45	15.7	48.4	50.23	49.7	17.44	15.92	19.86
DG ² CD-Net (TIES-Merging[67])	15.34	16.64	12.13	12.89	12.64	13.58	23.45	25.6	18.38	55.16	57.3	47.46	27.5	29.48	22.65
DG ² CD-Net [TA[21]]	15.17	16.52	11.8	12.78	12.58	13.29	23.21	25.69	17.34	55.16	57.31	48.87	26.76	27.91	23.96
DG ² CD-Net (Ours)	15.71	16.72	13.22	12.9	12.66	13.53	23.14	25.23	18.19	55.07	56.97	49.5	27.6	29.07	24.03
DG ² CD-Net* [Ours][LoRA[20]]	14.41	14.68	13.74	12.9	12.9	12.91	21.39	22.39	19.03	53.83	54.99	50.44	22.94	22.41	24.24

Table 11. Detailed comparison of our proposed DG²CD-Net on DG-GCD with respect to referred literature for DomainNet Dataset

K. Performance comparison with Domain Adaptation (DA) methods

Table 12 presents a performance comparison of our proposed method, DG²CD-Net, against two prominent methods, CROW (DA) and CDAD-Net (DA), across three benchmark datasets: PACS, Office-Home, and DomainNet. Notably, CDAD-Net (DA) represents a strong upper bound as it operates in the Domain Adaptation (DA) setting, leveraging access to target-domain data during training. In contrast, DG²CD-Net is designed for the more challenging Domain Generalization (DG) setting, where no target-domain information is available. While CROW (DA) is included in this supplementary comparison for reference, CDAD-Net (DA) serves as a more appropriate upper bound, given its superior performance.

Our method significantly outperforms CROW (DA) across all datasets in terms of overall accuracy, achieving a margin of +9.34% on PACS (73.30% vs. 63.96%), +3.39% on Office-Home (53.86% vs. 50.47%), and +2.81% on DomainNet (29.01% vs. 26.20%). Additionally, DG²CD-Net demonstrates robust generalization across both old and new classes, underscoring its adaptability in diverse scenarios. While CDAD-Net (DA) achieves higher performance due to its reliance on target domain data, the comparison highlights the inherent trade-off between the DA and DG settings. By including CROW (DA) results here, we provide a holistic view of baseline performance while emphasizing the relevance of CDAD-Net (DA) as the key upper bound in this context. This reinforces the practical value of DG²CD-Net in solving the domain generalization challenge without relying on target domain assumptions.

L. Additional Ablation Studies

Component Impact on Office-Home: Table 13 reveals the significant effects of essential components on the DG²CD-Net’s overall performance. Removing synthetic domains leads to a decrease of approximately 3.28% in the “All” metric,

Methods	Venue	Target-Domain	PACS			Office-Home			DomainNet		
			All	Old	New	All	Old	New	All	Old	New
CROW (DA) [60]	ICML'24	✓	63.96	61.78	68.65	50.47	54.50	39.71	26.20	26.60	25.80
DG²CD-Net (Ours)	–	×	73.30	75.28	72.56	53.86	53.37	54.33	29.01	30.38	25.46
CDAD-Net (DA) [40] [Upper bound]	CVPR-W'24	✓	83.25	87.58	77.35	67.55	72.42	63.44	70.28	76.46	65.19

Table 12. **Performance comparison** of CROW method with our method, as well as the upper bound CDAD-Net (DA), on all datasets.

underscoring its importance for generalization. The absence of episodic training results in a decrease of 2.33%, highlighting its role in model adaptability. The most considerable impact is observed with a fixed old/novel class split, which shows a reduction of 4.67% compared to the full model’s configuration. In contrast, the full implementation of DG²CD-Net achieves a comprehensive performance of 53.79% across all classes, demonstrating the effectiveness of dynamic weighting and the combined utility of all components in enhancing domain generalization and class discovery in the Office-Home dataset.

Model Variant	Office-Home		
	All	Old	New
✓ Without Synthetic Domain	50.51	50.58	50.31
✓ Without multi-global updates	51.46	50.77	52.26
✓ Static known/novel class split across episodes	49.11	56.23	38.85
Full DG²CD-Net (Proposed)	53.79	53.83	53.66

Table 13. Ablation study results on the impact of various components of DG²CD-Net for the Office-Home dataset.

Observation on Old-New class splits: Table 14 illustrates the impact of varying base (old) and novel (new) class splits on the performance of DG²CD-Net on the PACS dataset. We tested five different splits, ranging from 2 old classes with five novel classes to 6 old classes with one novel class. The results show that as the proportion of novel classes increases, the model’s performance on novel classes improves, but there is a slight decline in accuracy for base classes. This behavior highlights the challenge of maintaining a balance between recognizing old and novel categories. The best overall performance is observed with the 5-2 split, indicating that DG²CD-Net is more effective when the distribution of old and novel classes is moderately balanced.

Splits (Old-New)	PACS		
	All	Old	New
2 - 5	74.94	73.29	76.77
3 - 4	75.11	76.28	74.69
4 - 3	73.3	75.28	72.56
5 - 2	75.85	74.01	80.72
6 - 1	74.89	74.90	74.96

Table 14. Sensitivity on different Old-New class splits.

Effect of \mathcal{L}_{margin} : As shown in Figure 11, adding \mathcal{L}_{margin} improves accuracy across all categories. Old classes benefit the most with a 2.36% increase, indicating enhanced feature separation for well-known classes. New classes see a 2.23% improvement, suggesting better discrimination for novel classes. Overall, accuracy increases by 1.85%, demonstrating that \mathcal{L}_{margin} enhances class separability and generalization across both familiar and unseen categories.

Effect of m in \mathcal{L}_{margin} : Table 15 illustrates the average accuracy of our model across different margin values (m) in \mathcal{L}_{margin} . We experimented with different values of m ranging from 0.3 to 0.9. The highest scores were observed for $m = 0.7$, indicating that this margin setting provides better separation between known and novel classes.

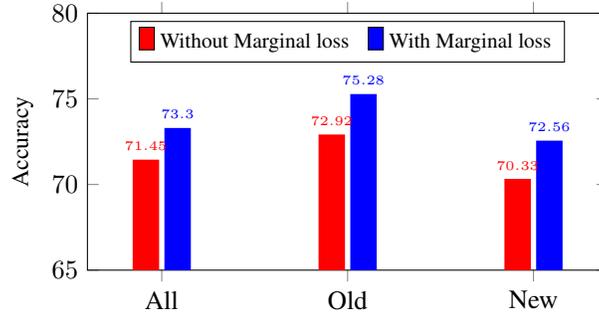


Figure 11. Accuracy comparison of Old, All, and New categories with and without $\mathcal{L}_{\text{margin}}$ loss.

m	PACS		
	All	Old	New
0.3	71.50	74.48	69.05
0.4	71.27	75.58	67.92
0.5	69.82	74.00	66.23
0.6	63.13	65.42	61.99
0.7	73.30	75.28	72.56
0.8	72.24	75.77	69.92
0.9	71.91	75.69	68.98

Table 15. Average accuracy for different sensitivity of the hyper-parameter m .

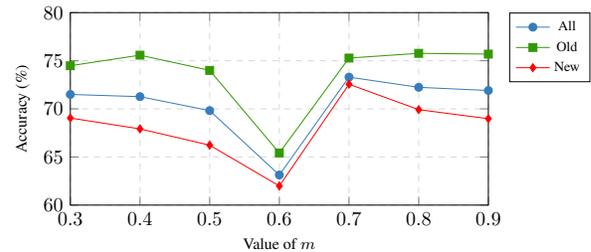


Figure 12. The relationship between margin m and accuracy (All, Old, New). The margin m in the loss function influences class separation, with peak accuracy observed at $m = 0.7$.

M. Effect of initialization of backbone

Backbone selection significantly impacts the performance of DG²CD-Net in domain generalization. Initially, we employed ViT-B/16 with DINO initialization for comparability with prior work, but inspired by DINO-v2’s advancements, we expanded our study to include ViT-B/16 with DINO-v2 initialization. Additionally, we have experimented with ResNet-50 with CLIP and ImageNet initializations and ViT-B/16 with CLIP initialization. Table 16 summarizes the results on PACS.

Model	Backbone	PACS		
		All	Old	New
ResNet-50	CLIP [39]	25.39	20.97	29.33
ResNet-50	ImageNet [18]	54.98	64.18	45.33
ViT-B/16	DINO [6]	73.30	75.28	72.56
ViT-B/16	DINO-v2 [33]	87.71	90.67	84.91
ViT-B/16	CLIP [39]	90.07	92.25	87.72

Table 16. Performance Comparison of DG²CD-Net with different backbones on the PACS Dataset

The results highlight the superiority of the DINO-v2 and CLIP-based models, with CLIP achieving the highest performance. The strong results of ViT-B/16 (CLIP) suggest that pre-training with vision-language data improves generalization. Given these findings, we recommend DINO-v2 and CLIP-based ViT as strong baselines for future domain generalization studies.

N. Effect of LoRA Fine-Tuning

In our experiments, we incorporated LoRA (Low-Rank Adaptation) into DG² CD-Net to improve the efficiency of fine-tuning while minimizing memory overhead. Unlike full fine-tuning, LoRA updates a small subset of parameters while keeping pre-trained weights frozen, thereby reducing catastrophic forgetting and enhancing adaptation and generalization.

Method	Trainable Parameters (K)	Total Parameters (K)	Percentage (%)
DG ² CD-Net (Vanilla)	7,088	85,799	8.261
DG ² CD-Net* (LoRA)	98	85,799	0.115

Table 17. Comparison of model parameters *with* and *without* LoRA fine-tuning.

Table 18 presents a performance comparison of DG²CD-Net using different LoRA-based adapters, including *LoRA* [20], *DoRA* [29], and *AdaLoRA* [71]. These adapters aim to improve the model efficiency while maintaining high accuracy, with LoRA achieving the best performance.

Adapters	All (%)
LoRA [20]	75.21
DoRA [29]	74.11
AdaLoRA [71]	74.20

Table 18. Performance comparison of DG²CD-Net with different LoRA Adapters

These findings emphasize the efficiency of LoRA in reducing the computational burden of fine-tuning while maintaining the model’s overall capacity. Given the benefits, we recommend adopting LoRA-based fine-tuning in future research for domain generalization tasks to optimize memory usage and training speed without sacrificing performance.

O. Limitations and Future Work

While DG²CD-Net demonstrates strong performance in domain generalization and category discovery, there are areas for future enhancement. One aspect that can be improved is the reliance on synthetic domain generation, which, while effective, can be optimized to reduce computational costs. Exploring more streamlined approaches to synthetic domain creation or alternative techniques that do not require synthetic data could further improve scalability and efficiency. Additionally, the episodic training framework, though beneficial for adaptation, demands considerable computational resources, especially when applied to large-scale datasets like DomainNet. Optimizing this process could make the method more feasible for real-world, large-scale applications.

In future work, efforts can focus on enhancing the efficiency of both synthetic domain generation and the episodic training process. Advanced techniques for model merging can also be explored to further improve performance. Addressing challenges like data imbalance, which is common in real-world scenarios, will strengthen the model’s robustness and adaptability. Overall, extending DG²CD-Net in these directions holds great promise for developing even more scalable and effective solutions for complex tasks in domain generalization.

References

- [1] Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. *Advances in Neural Information Processing Systems*, 35:8265–8277, 2022. 3
- [2] Joachim A Behar, Jeremy Levy, and Leo Anthony Celi. Generalization in medical ai: a perspective on developing scalable models. *arXiv preprint arXiv:2311.05418*, 2023. 1
- [3] Prathmesh Bele, Valay Bundele, Avigyan Bhattacharya, Ankit Jha, Gemma Roig, and Biplab Banerjee. Learning class and domain augmentations for single-source open-domain generalization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1816–1826, 2024. 1, 2, 3, 6
- [4] Shirsha Bose, Ankit Jha, Hitesh Kandala, and Biplab Banerjee. Beyond boundaries: A novel data-augmentation discourse for open domain generalization. *Transactions on Machine Learning Research*, 2023. 3
- [5] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 2, 4
- [6] 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. 6, 18
- [7] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34:22405–22418, 2021. 3
- [8] Junyi Chai, Hao Zeng, Anming Li, and Eric WT Ngai. Deep learning in computer vision: A critical review of emerging techniques and application scenarios. *Machine Learning with Applications*, 6:100134, 2021. 1
- [9] Florent Chironi, 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. 3, 10
- [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. 6, 7, 10
- [11] Ilke Cugu, Massimiliano Mancini, Yanbei Chen, and Zeynep Akata. Attention consistency on visual corruptions for single-source domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4165–4174, 2022. 3
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 3, 6, 15, 16
- [13] Thomas Eiter and Heikki Mannila. Computing discrete fréchet distance. 1994. 4
- [14] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. 3
- [15] Kai Han, Andrea Vedaldi, and Andrew Zisserman. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8401–8409, 2019. 1, 10
- [16] Kai Han, Sylvestre-Alvise Rebuffi, Sebastien Ehrhardt, Andrea Vedaldi, and Andrew Zisserman. Automatically discovering and learning new visual categories with ranking statistics. In *International Conference on Learning Representations*, 2020. 3, 10
- [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*, 44(10):6767–6781, 2021. 3, 10
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 18
- [19] Yen-Chang Hsu, Zhaoyang Lv, and Zsolt Kira. Learning to cluster in order to transfer across domains and tasks. In *International Conference on Learning Representations*, 2018. 1, 3, 10
- [20] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 6, 7, 15, 16, 19
- [21] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 3, 4, 5, 6, 7, 8, 15, 16
- [22] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020. 1
- [23] Xisen Jin, Xiang Ren, Daniel Preotiuc-Pietro, and Pengxiang Cheng. Dataless knowledge fusion by merging weights of language models. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [24] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 3, 6
- [25] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 6, 10

- [26] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 3
- [27] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 3
- [28] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zihua Zhang. On the convergence of fedavg on non-iid data. In *International Conference on Learning Representations*. 3
- [29] Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. DoRA: Weight-decomposed low-rank adaptation. In *Proceedings of the 41st International Conference on Machine Learning*, pages 32100–32121. PMLR, 2024. 19
- [30] Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances in Neural Information Processing Systems*, 35:17703–17716, 2022. 2, 8
- [31] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International conference on machine learning*, pages 10–18. PMLR, 2013. 3
- [32] OpenAI. Chatgpt, 2024. <https://chat.openai.com/>. 4
- [33] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. Featured Certification. 18
- [34] Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. *Advances in Neural Information Processing Systems*, 36, 2023. 3
- [35] Kunyu Peng, Di Wen, Kailun Yang, Ao Luo, Yufan Chen, Jia Fu, M Saquib Sarfraz, Alina Roitberg, and Rainer Stiefelhagen. Advancing open-set domain generalization using evidential bi-level hardest domain scheduler. *Advances in Neural Information Processing Systems*, 37:85412–85440, 2025. 3
- [36] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 6, 10
- [37] Xi Peng, Fengchun Qiao, and Long Zhao. Out-of-domain generalization from a single source: An uncertainty quantification approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(3):1775–1787, 2022. 3
- [38] Nan Pu, Zhun Zhong, and Nicu Sebe. Dynamic conceptional contrastive learning for generalized category discovery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7579–7588, 2023. 1, 3, 10
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 18
- [40] Sai Bhargav Rongali, Sarthak Mehrotra, Ankit Jha, Shirsha Bose, Tanisha Gupta, Mainak Singha, Biplab Banerjee, et al. Cdad-net: Bridging domain gaps in generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2616–2626, 2024. 1, 3, 5, 6, 7, 8, 10, 15, 16, 17
- [41] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 153–168, 2018. 1, 2, 5, 12
- [42] Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9624–9633, 2021. 3
- [43] Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9624–9633, 2021. 6, 7
- [44] Mainak Singha, Ankit Jha, Shirsha Bose, Ashwin Nair, Moloud Abdar, and Biplab Banerjee. Unknown prompt the only lacuna: Unveiling clip’s potential for open domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13309–13319, 2024. 3
- [45] Tao et al. Sun. Shift: a synthetic driving dataset for continuous multi-task domain adaptation. In *CVPR*, 2022. 1
- [46] Yi-Lin Sung, Linjie Li, Kevin Lin, Zhe Gan, Mohit Bansal, and Lijuan Wang. An empirical study of multimodal model merging. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023. 3
- [47] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 7
- [48] Jesper E Van Engelen and Holger H Hoos. A survey on semi-supervised learning. *Machine learning*, 109(2):373–440, 2020. 1
- [49] 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, 4, 6, 7, 10, 15, 16
- [50] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017. 4, 6, 10

- [51] Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial intelligence review*, 18:77–95, 2002. 5
- [52] Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. Deep learning for computer vision: A brief review. *Computational intelligence and neuroscience*, 2018(1):7068349, 2018. 1
- [53] Hongjun Wang, Sagar Vaze, and Kai Han. Hilo: A learning framework for generalized category discovery robust to domain shifts, 2024. 3
- [54] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and S Yu Philip. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8): 8052–8072, 2022. 2
- [55] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018. 1
- [56] Xiran Wang, Jian Zhang, Lei Qi, and Yinghuan Shi. Generalizable decision boundaries: Dualistic meta-learning for open set domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11564–11573, 2023. 3
- [57] Xiran Wang, Jian Zhang, Lei Qi, and Yinghuan Shi. Generalizable decision boundaries: Dualistic meta-learning for open set domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11564–11573, 2023. 3
- [58] Yicheng Wang, Feng Liu, Junmin Liu, Zhen Fang, and Kai Sun. Exclusive style removal for cross domain novel class discovery, 2024. 3
- [59] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh. Learning to diversify for single domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 834–843, 2021. 3
- [60] Shuo Wen and Maria Brbic. Cross-domain open-world discovery. *arXiv preprint arXiv:2406.11422*, 2024. 3, 7, 17
- [61] 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. 3, 6, 10, 15, 16
- [62] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR, 2022. 3
- [63] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hananeh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7959–7971, 2022. 3
- [64] Qinwei Xu, Ruipeng Zhang, Yi-Yan Wu, Ya Zhang, Ning Liu, and Yanfeng Wang. Simde: A simple domain expansion approach for single-source domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4798–4808, 2023. 3
- [65] Prateek Yadav and Mohit Bansal. Exclusive supermask subnetwork training for continual learning. *arXiv preprint arXiv:2210.10209*, 2022. 3
- [66] Prateek Yadav, Qing Sun, Hantian Ding, Xiaopeng Li, Dejiao Zhang, Ming Tan, Parminder Bhatia, Xiaofei Ma, Ramesh Nallapati, Murali Krishna Ramanathan, et al. Exploring continual learning for code generation models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 782–792, 2023. 3
- [67] Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 3, 6, 7, 15, 16
- [68] Jee Seok Yoon, Kwansoek Oh, Yooseung Shin, Maciej A Mazurowski, and Heung-Il Suk. Domain generalization for medical image analysis: A survey. *arXiv preprint arXiv:2310.08598*, 2023. 1
- [69] Xiaohua Zhai, Avital Oliver, Alexander Kolesnikov, and Lucas Beyer. S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1476–1485, 2019. 1
- [70] Ling Zhang, Xiaosong Wang, Dong Yang, Thomas Sanford, Stephanie Harmon, Baris Turkbey, Holger Roth, Andriy Myronenko, Daguang Xu, and Ziyue Xu. When unseen domain generalization is unnecessary? rethinking data augmentation. *arXiv preprint arXiv:1906.03347*, 2019. 1
- [71] Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive budget allocation for parameter-efficient fine-tuning. In *The Eleventh International Conference on Learning Representations*, 2023. 19
- [72] Sheng Zhang, Salman Khan, Zhiqiang Shen, Muzammal Naseer, Guangyi Chen, and Fahad Shahbaz Khan. Promptcal: Contrastive affinity learning via auxiliary prompts for generalized novel category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3479–3488, 2023. 1, 3, 10
- [73] Bingchen Zhao and Kai Han. Novel visual category discovery with dual ranking statistics and mutual knowledge distillation. *Advances in Neural Information Processing Systems*, 34:22982–22994, 2021. 10
- [74] 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. 10
- [75] Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10867–10875, 2021. 3, 10

- [76] Zhun Zhong, Linchao Zhu, Zhiming Luo, Shaozi Li, Yi Yang, and Nicu Sebe. Openmix: Reviving known knowledge for discovering novel visual categories in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9462–9470, 2021. [3](#), [10](#)
- [77] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 561–578. Springer, 2020. [3](#)
- [78] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021. [3](#)
- [79] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4396–4415, 2022. [1](#), [2](#), [6](#)
- [80] Ronghang Zhu and Sheng Li. Crossmatch: Cross-classifier consistency regularization for open-set single domain generalization. In *International conference on learning representations*, 2022. [3](#)