

# Learning Fine-grained Domain Generalization via Hyperbolic State Space Hallucination

Qi Bi<sup>1</sup>, Jingjun Yi<sup>1</sup>, Haolan Zhan<sup>2</sup>, Wei Ji<sup>3\*</sup>, Gui-Song Xia<sup>1\*</sup>

<sup>1</sup>School of Artificial Intelligence, Wuhan University, Wuhan, China

<sup>2</sup>Faculty of Information Technology, Monash University, Melbourne, Australia

<sup>3</sup>School of Medicine, Yale University, New Haven, United States

{q.bi, guisong.xia}@whu.edu.cn, wei.ji@yale.edu

## Abstract

Fine-grained domain generalization (FGDG) aims to learn a fine-grained representation that can be well generalized to unseen target domains when only trained on the source domain data. Compared with generic domain generalization, FGDG is particularly challenging in that the fine-grained category can be only discerned by some subtle and tiny patterns. Such patterns are particularly fragile under the cross-domain style shifts caused by illumination, color and etc. To push this frontier, this paper presents a novel Hyperbolic State Space Hallucination (HSSH) method. It consists of two key components, namely, state space hallucination (SSH) and hyperbolic manifold consistency (HMC). SSH enriches the style diversity for the state embeddings by firstly extrapolating and then hallucinating the source images. Then, the pre- and post- style hallucinate state embeddings are projected into the hyperbolic manifold. The hyperbolic state space models the high-order statistics, and allows a better discernment of the fine-grained patterns. Finally, the hyperbolic distance is minimized, so that the impact of style variation on fine-grained patterns can be eliminated. Experiments on three FGDG benchmarks demonstrate its state-of-the-art performance.

**Code** — <https://github.com/BiQiWHU/HSSH>

## Introduction

Visual domain generalization is a fundamental task in both computer vision and machine learning, which handles the scenarios when the training samples and testing samples are not independently and identically distributed (i.i.d.). It has been extensively studied in the past decade, which usually assumes that the domain gap is caused by the style variation, while the content (i.e., semantic category) is stable (Yan et al. 2020; Chen et al. 2023; Wang et al. 2023). Great success has been witnessed in advancing the generalization ability of a deep learning model in a variety of applications, such as semantic segmentation and object detection (Bi, You, and Gevers 2024a,b; Zhao et al. 2022; Lee et al. 2022; Yi et al. 2024b; Bi et al. 2024b). However, learning a domain generalized model for fine-grained visual categorization, termed as *Fine-grained domain generalization* (FGDG), remains less explored.

\*Corresponding authors

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

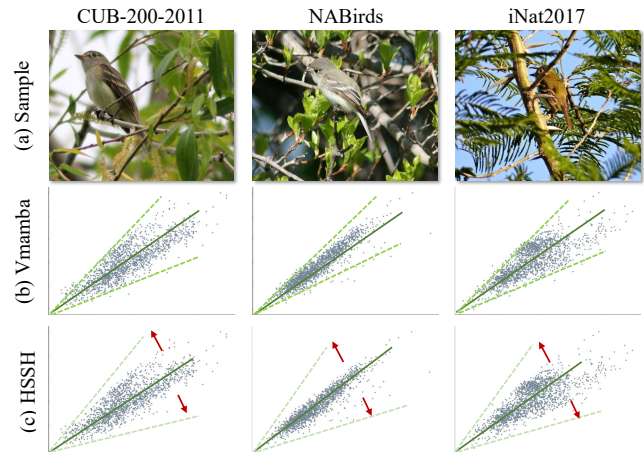


Figure 1: In each fine-grained domain (first row), due to the inter-category similarity, the embedding style from the VMamba baseline (second row) is limited. The proposed HSSH (third row) enriches the style diversity of the state embedding, which benefits the robustness to unseen domains. The x- and y- axis refers to the mean and standard deviation value in the style space, ranging from 0 to 1.

FGDG holds both technical and practical significance. From a technical view, in generic domain generalization, the cross-domain content (i.e., semantic category) is stable and easy to differentiate. However, in FGDG, the fine-grained categories from a subordinate coarse-grained category can only be distinguished from some subtle and tiny patterns. This problem becomes even more challenging when the cross-domain styles shift. The shallower features of a deep model are sensitive to the style shift, posed by varied illumination, color, structure, etc. The impact of such permutation accumulates from shallower to deeper features, making it more difficult to perceive the subtle fine-grained patterns. From a practical view, FGDG also deserves further study. Existing fine-grained visual categorization (FGVC) datasets are still small in amount (10K). A pre-trained FGVC model is likely to encounter unseen and out-of-domain samples when doing real-world inference.

The emerging selective state space model (SSM) (Gu et al. 2021; Gu, Goel, and Re 2022), exemplified by the



Figure 2: Hyperbolic manifold provides a feasible path to learn hierarchical and high-order statistics, which has potential to distinguish a fine-grained category from other fine-grained categories under the same coarse-grained category.

VMamba (Zhu et al. 2024) and Vision Mamba (Liu et al. 2024), has become a recent paradigm. It is able to exploit the relation between local- and global- representations by selective scan mechanism and recurrent modeling from the patch embeddings of an image. This property makes SSM have great potential for FGDG, as the local subtle patterns can be handled by the selective scan mechanism and patch embedding, while the scene content can be handled by the recurrent modeling to exploit the long-range dependency.

In this paper, we aim to advance SSM for FGDG, by allowing it to simultaneously improve the robustness to cross-domain style shift and exploit the fine-grained semantic category representation. To realize the first objective, we propose a state space hallucination (SSH) technique, which hallucinates the state embeddings from each VMamba layer. By implementing linear extrapolation on the per-batch state embeddings, it enriches the style diversity than the VMamba baseline (shown in Fig. 1). The enriched styles allow the state embeddings to be less sensitive to the domain shift when inference on unseen target domains, benefiting the learning of subtle fine-grained patterns.

To realize the second objective, a hyperbolic manifold consistency (HMC) learning scheme is proposed. Specifically, both pre- and post- hallucinated state embeddings are projected into the hyperbolic manifold, which is well recognized to preserve the hierarchical and high-order statistics (Ganea, Bécigneul, and Hofmann 2018; Yu and De Sa 2019). The hyperbolic manifold allows the state embeddings from the same fine-grained semantic category to be as close as possible, while the state embeddings from different fine-grained categories to be as apart as possible (shown in Fig. 2). Afterwards, the hyperbolic distance between the

pre- and post- hallucinated state embedding is minimized, so as to alleviate the impact from cross-domain styles on differentiating the fine-grained semantic category.

The proposed hyperbolic state space hallucination (HSSH) method is empowered by the above two key components, namely, SSH and HMC. Notably, it not only makes an early exploration of how to harness SSM for FGDG, but also of how to embed the hyperbolic manifold into the state space. The proposed HSSH is validated on three FGDG benchmarks. In each FGDG benchmark, two or more fine-grained visual categorization (FGVC) datasets from the same coarse-grained category (e.g., bird) are used as the source and target domain, and the common fine-grained categories among these datasets are selected.

Our contributions can be summarized as follows.

- We initiate an early exploration to harness State Space Models (SSM) for fine-grained domain generalization.
- A Hyperbolic State Space Hallucination (HSSH) method is proposed, which innovatively embeds the hyperbolic manifold into the selective state space.
- We propose a state space hallucination (SSH) component to enrich the style diversity for SSM, and a hyperbolic manifold consistency (HMC) component to learn discriminative fine-grained categories under domain shift.
- Extensive experiments show that the proposed HSSH significantly outperforms the state-of-the-art by up to 16.14%, 26.01% and 12.54% on CUB-Paintings, RS-FGDG and Birds-31, respectively.

## Related Work

**Fine-grained Domain Generalization** Over the past decade, Fine-grained Visual Categorization (FGVC) have been extensively studied (Zhao et al. 2021; Hu et al. 2021; Du et al. 2021; Yang et al. 2022). However, these methods assume that the training and testing images are i.i.d., which is far from the reality. FGDG handles this challenge. Specifically, a progressive learning scheme is proposed to maintain the consistency of fine-grained semantics for FGVC on different domains (Du et al. 2021). A progressive adversarial network (PAN) along with a curriculum-based adversarial learning scheme is proposed to align fine-grained categories across domains (Wang et al. 2021). More recently, a feature structuralization method is proposed to stabilize the fine-grained semantics across domains (Yu et al. 2024).

**Style Hallucination** has been recognized as a common solution to enrich style diversity (Huang and Belongie 2017; Chen et al. 2021; Zhang et al. 2022). It has also been widely explored in domain generalized semantic segmentation (Zhao et al. 2022; Zhong et al. 2022; Zhao et al. 2024; Yi et al. 2024a), where the augmented styles enhance the generalization to unseen domains that have different styles.

**State Space Models (SSM)** (Kalman 1960) is fundamental for various fields, and has recently shown great success for deep representation learning (Gu et al. 2021; Gu, Goel, and Re 2022; Gu and Dao 2023; Bi et al. 2024a). This has led to the development of Mamba variants such as Vision Mamba (Liu et al. 2024) and VMamba (Zhu et al. 2024) for computer vision. Nevertheless, the generalization ability of

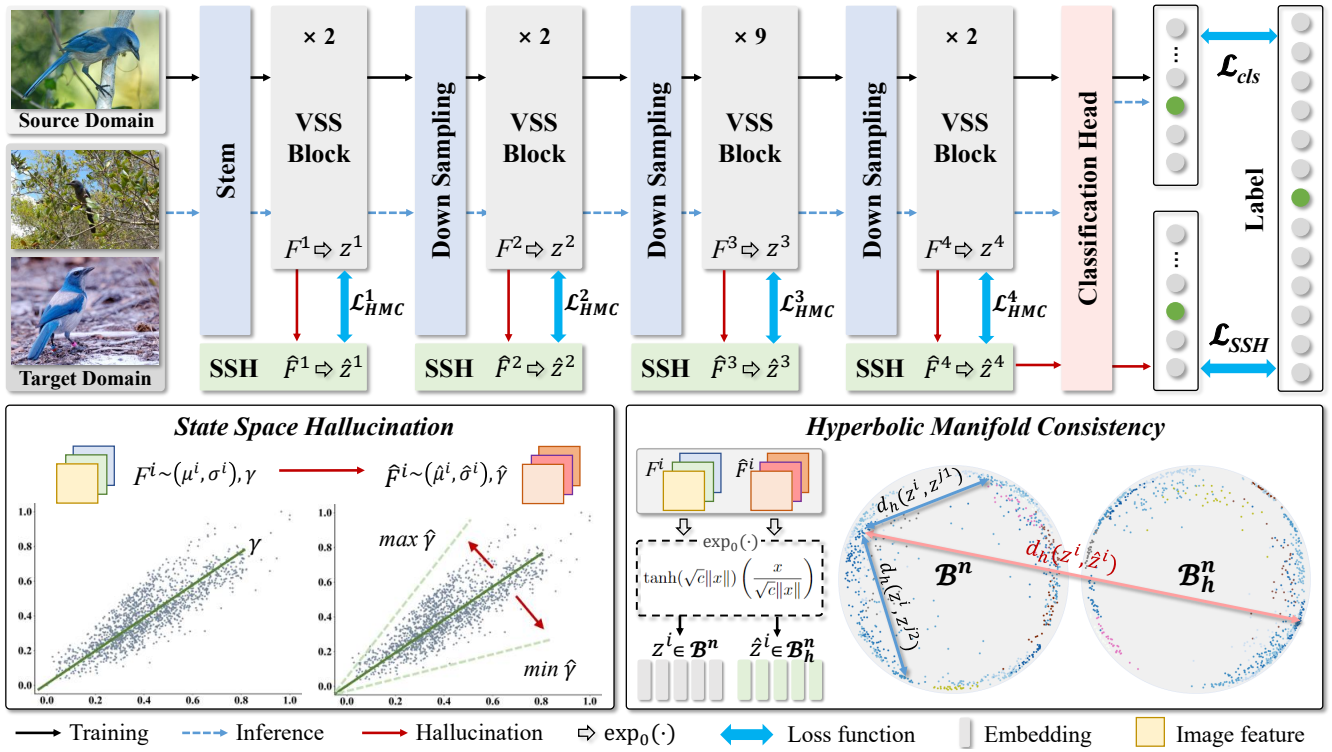


Figure 3: The proposed hyperbolic state space hallucination (HSSH) is empowered by two key components, namely, state space hallucination (SSH) and hyperbolic manifold consistency (HMC). SSH enriches the style diversity for the state embeddings by firstly extrapolating and then hallucinating from the source images. HMC minimizes the hyperbolic distance between the pre- and post- hallucinated samples, so that the impact from the style variation on fine-grained patterns can be eliminated.

SSM on unseen target domains, especially for fine-grained semantic categories, remains to be explored.

**Hyperbolic Manifold Learning** has garnered significant interest (Ganea, Bécigneul, and Hofmann 2018; Yu and De Sa 2019) and have been applied to many computer vision tasks (Yang et al. 2023; Desai et al. 2023; Hu et al. 2024; Kong et al. 2024; Mettes et al. 2024; Xu et al. 2023; Weber et al. 2020; Desai et al. 2023; Hu et al. 2024; Kong et al. 2024). However, how to embed hyperbolic alignment within state space or for FGDG remains an open question.

## Methodology

### Preliminaries & Overview

Given  $K$  unseen target domains  $\mathcal{D}_1, \dots, \mathcal{D}_K$  and a single source domain  $\mathcal{D}_{K+1}$ , we denote the joint image and label set within domain  $\mathcal{D}_k$  as  $\{(x_n^k, y_{c,n}^k, y_{f,n}^k)\}_{n=1}^{N_k}$ , where  $k = 1, 2, \dots, K$ , and  $N_k$  represents the number of samples in domain  $\mathcal{D}_k$ . Here,  $y_{c,n}^k$  and  $y_{f,n}^k$  denote the coarse-grained and fine-grained label of the sample  $x_n^k$ , respectively. Each sample among these domains shares a common coarse-grained category  $y_{c,n}^k$  but varies in different fine-grained categories  $y_{f,n}^k$ . Our objective is to learn a model  $f_\theta : x_n^k \rightarrow y_{f,n}^k$ , trained solely on the single source domain  $\mathcal{D}_{K+1}$ , that generalizes to the unseen target domains.

Fig. 3 gives an overview of the proposed Hyperbolic

State Space Hallucination (HSSH) scheme. On top of the VMamba backbone, it consists of two key components, namely, State Space Hallucination (SSH), and Hyperbolic Manifold Consistency (HMC). SSH enriches the style space  $\mathbb{S}$  and then hallucinates the state embeddings of samples  $x_n^{k+1}$  from the source domain  $\mathcal{D}_{K+1}$ . HMC projects the pre- and post- hallucinated state embeddings into the hyperbolic manifolds, where they are aligned by minimizing the geodesic distance  $d$  between hyperbolic embeddings  $z$ . It helps promote stable representation of the fine-grained semantic separability in the state space.

### State Space Hallucination

For FGVC, the key challenge is to identify the subtle and tiny fine-grained patterns that represent the fine-grained semantics from numerous patches of an image. The Mamba architecture (Gu and Dao 2023) encodes the input data into sequences, implements the selective scan mechanism and recurrent modeling (Gu, Goel, and Re 2022) to highlight the relevant information (Lu et al. 2024). This property has great potential to discern the key patterns, but the robustness to style variation needs to be enhanced for cross-domain scenario. The state space hallucination (SSH) component is proposed to push this frontier.

The Vision Mamba encoder (Zhu et al. 2024) consists of four blocks from shallow to deep. Assume the state embeddings from the  $i$ -th block are denoted as

$\mathbf{F}^i \in \mathbb{R}^{B \times C_i \times H_i \times W_i}$ , where  $C_i, H_i, W_i$  denote the channel, height and width of a patch, respectively. Following the existing definition (Huang and Belongie 2017; Chen et al. 2021; Zhang et al. 2022), the style of the state embeddings is quantified by computing the channel-wise mean  $\mu$  and standard deviation  $\sigma$ . Specifically, given a batch of samples from the source domain  $\{x_n\}_{n=1}^B$ , where  $B$  denotes the batch size, the channel-wise mean  $\mu^i \in \mathbb{R}^{B \times C_i}$  and deviation  $\sigma^i \in \mathbb{R}^{B \times C_i}$  can be mathematically computed as:

$$\begin{aligned} \mu^i &= \frac{1}{H_i W_i} \mathbf{F}^i, \\ \sigma^i &= \sqrt{\frac{1}{H_i W_i} \sum_{h,w \in H_i, W_i} (\mathbf{F}^i - \mu^i)(\mathbf{F}^i - \mu^i)^T}. \end{aligned} \quad (1)$$

As shown in Fig. 1b, the style distribution of the state embeddings is highly constrained in the style space  $\mathcal{S} \subset \mathbb{R}^{[0,1] \times [0,1]}$ , where the upper and lower slopes (green lines) are relatively close to each other. Such limited diversity makes it less generalized to unseen target domains with different styles. To handle this, we enrich its diversity by style hallucination, which enlarges the differences between the upper and lower slopes (in Fig. 1c). The general idea is to introduce another slope in  $\mathcal{S}$  to represent the hallucinated styles without prior knowledge of other domains. Considering  $(\mu_{j,k}^i, \sigma_{j,k}^i)$  as samples points, where  $0 \leq j \leq B, 0 \leq k \leq C_i$ , we approximate the slope value  $\gamma^i$  from the samples in a batch  $\{x_n\}_{n=1}^B$  by the least square method:

$$\gamma^i = \frac{\sum_{j=1}^B \sum_{k=1}^{C_i} (\mu_{j,k}^i - \bar{\mu}^i)(\sigma_{j,k}^i - \bar{\sigma}^i)}{\sum_{j=1}^B \sum_{k=1}^{C_i} (\mu_{j,k}^i - \bar{\mu}^i)^2}, \quad (2)$$

where  $\bar{\mu}^i$  and  $\bar{\sigma}^i$  denote the average values of  $\mu^i$  and  $\sigma^i$ .

Then, the style range is expanded by altering the slope value and sampling style points on the line of the new slope. We calculate the maximum slope  $\max \gamma^i$  and the minimum slope  $\min \gamma^i$  of the original state embeddings. Then, we sample new points on the expanded style regions with a broader range of slope values ( $\min \tilde{\gamma} \leq \tilde{\gamma} \leq \max \tilde{\gamma}$ ) by  $\min \tilde{\gamma} = 2 \min \gamma - \max \gamma$  and  $\max \tilde{\gamma} = 2 \max \gamma - \min \gamma$ . Here, the new style slopes are sampled with random slope  $\tilde{\gamma} \sim [\min \tilde{\gamma}, \max \tilde{\gamma}]$ . Then we obtain hallucinated features  $\tilde{\mathbf{F}}^i$  from sampled points  $(\tilde{\mu}^i, \tilde{\sigma}^i)$ , given by:

$$\tilde{\mathbf{F}}^i = \tilde{\sigma}^i \frac{\mathbf{F}^i - \mu^i}{\sigma^i} + \tilde{\mu}^i. \quad (3)$$

### Hyperbolic Manifold Consistency

The enrichment of style diversity does not necessarily ensure that the state space can precisely discern the subtle and tiny fine-grained patterns from a number of image patches. To address this issue, we further project the state embeddings into the hyperbolic manifold, which is well recognized to represent the hierarchical and high-order information (Ganea, Bécigneul, and Hofmann 2018; Yu and De Sa 2019). This property allows the samples from different fine-grained categories to fall apart. We propose Hyperbolic Manifold Consistency (HMC) to constrain the pre- and post-

hallucinated style embeddings, so that the fine-grained representation can be consistent despite the style shift.

Specifically, let the hallucinated features from the four encoder blocks are denoted as  $\tilde{\mathbf{F}}^1, \tilde{\mathbf{F}}^2, \tilde{\mathbf{F}}^3$  and  $\tilde{\mathbf{F}}^4$ , respectively. On the other hand, Poincaré ball is selected to as the consistency manifold since it has been widely utilized in computer vision community. The Poincaré ball model is characterized by the pair  $(\mathcal{B}_c^n, g_c^B)$ , where the manifold is defined as  $\mathcal{B}_c^n = \{x \in \mathbb{R}^n : c\|x\|^2 < 1\}$ . As the fine-grained patterns exist in the Euclidean space, a transformation is needed to work within the hyperbolic space of the Poincaré ball. This objective is achieved by mapping the state embedding  $\mathbf{F}^i$  using the exponential map centered at a reference point  $v$ :

$$\exp_v^c(\mathbf{F}^i) = v \oplus_c \left( \tanh \left( \frac{\sqrt{c}\|\mathbf{F}^i\|}{2} \right) \frac{\mathbf{F}^i}{\sqrt{c}\|\mathbf{F}^i\|} \right), \quad (4)$$

where  $\oplus_c$  denotes the Möbius addition, given by

$$v \oplus_c w = \frac{(1 + 2c\langle v, w \rangle + c\|w\|^2)v + (1 - c\|v\|^2)w}{1 + 2c\langle v, w \rangle + c^2\|v\|^2\|w\|^2}. \quad (5)$$

Practically, the point  $v$  is often taken as the origin. The exponential map expression can be simplified as

$$\exp_0(\mathbf{F}^i) = \tanh(\sqrt{c}\|\mathbf{F}^i\|) \left( \frac{\mathbf{F}^i}{\sqrt{c}\|\mathbf{F}^i\|} \right). \quad (6)$$

Let  $z^i \in \mathbb{R}^{B \times C \times H_i \times W_i}, \tilde{z}^i \in \mathbb{R}^{B \times C \times H_i \times W_i}$  denote the hyperbolic embedding of the pre-hallucinated state embedding  $\mathbf{F}^i$  and the post-hallucinated state embedding  $\tilde{\mathbf{F}}^i$ , the projection to the Poincaré hyperbole (by Eq. 6) can be reformulated as  $z^i = \exp_0(\mathbf{F}^i), \tilde{z}^i = \exp_0(\tilde{\mathbf{F}}^i)$ . Then, the hyperbolic distance  $d_h$  can be calculated as

$$d_h(z^i, \tilde{z}^i) = \frac{2}{\sqrt{c}} \tanh^{-1} \left( \sqrt{c} \left\| -z^i \oplus_c \tilde{z}^i \right\| \right), \quad (7)$$

where  $c$  is a hyperparameter that controls the radius of the hyperbolic space, and by default it is set to 0.1. To ensure that the fine-grained representation is consistent between the pre- and post- hallucinated hyperbolic embeddings, we devise a hyperbolic manifold consistency loss:

$$\mathcal{L}_{HMC} = \left( d_h(z^i, \tilde{z}^i) + \frac{1}{N_S} \sum_{(i,j) \in S} d_h(z^i, z^j) \right), \quad (8)$$

where  $S$  denotes the set of sample pairs  $(z_i, z_j)$  that belong to the same coarse-grained category, and  $N_S$  denotes the number of these pairs.

### Implementation Details

The pre- and post- hallucinated state embeddings  $\mathbf{F}_c^4$  and  $\tilde{\mathbf{F}}_c^4$  from the last block of the VMamba encoder are fed into the classifier  $\phi(\cdot)$ , optimized by the classification loss function  $\mathcal{L}_{cls}$  and  $\tilde{\mathcal{L}}_{cls}$ , respectively.  $\phi(\cdot)$  is a linear layer with an input/output dimension of 768/category number, respectively. Assume that  $C$  represents the category number. Then, the classification loss is defined as

$$\mathcal{L}_{cls} = -\frac{1}{B} \sum_{i=1}^B \sum_{j=1}^C y_{i,j} \log(\phi(\mathbf{F}^4)_{i,j}). \quad (9)$$

The total loss  $\mathcal{L}$  is a linear combination between the classification loss and the hyperbolic manifold consistency  $\mathcal{L}_{HMC}$ , given by

$$\mathcal{L} = \mathcal{L}_{cls} + \tilde{\mathcal{L}}_{cls} + \lambda \mathcal{L}_{HMC}, \quad (10)$$

where  $\lambda$  is a hyper-parameter used to balance semantic consistency in the state space and hierarchical consistency in the hyperbolic space. Empirically,  $\lambda$  is set to 0.5.

For all experiments across the three FGDG settings, the Adam optimizer is employed with a learning rate of  $1 \times 10^{-4}$ , and momentum parameters set to 0.9 and 0.99. The training process spans 100 epochs.

## Experiments

### Dataset and Evaluation Protocols

**CUB-Paintings** encompasses two domains: Caltech-UCSD Birds-200-2011 (CUB-200-2011, denoted as C) (Wah et al. 2011) and CUB-200-Paintings (CUB-P, denoted as P) (Wang et al. 2020). CUB-200-2011 contains 11,788 images across 200 bird subcategories. CUB-P shares the same categorization system and includes 3,047 images featuring artistic representations such as watercolors, oil paintings, pencil drawings, stamps, and cartoons. These datasets exhibit notable domain shifts due to different styles. Furthermore, both domains can be grouped into 38 coarse categories.

**RS-FGDG** comprises two domains: Million-AID (Long et al. 2021) (MAID, denoted as M) and NWPU-RESISC45 (Cheng, Han, and Lu 2017) (NWPU, denoted as N). Million-AID contains 1,000,848 scenes organized into 51 fine-grained categories. NWPU includes 31,500 images across 45 scene classes, each with 700 images. For domain generalization, the 22 common fine-grained categories from 8 coarse-grained categories are used.

**Birds-31** comprises three domains: Caltech-UCSD Birds-200-2011 (CUB-200-2011, denoted as C) (Wah et al. 2011), NABirds (denoted as N) (Van Horn et al. 2015), and iNaturalist2017 (iNat2017, denoted as I) (Van Horn et al. 2018). NABirds consists of 555 categories, with a total of 48,562 images. iNat2017 contains 859,000 images spanning over 5,000 species of plants and animals. For FGDG, 31 common categories were selected, resulting in a total of 7,693 images, with 1,848 from CUB-200-2011, 2,988 from NABirds, and 2,857 from iNat2017. These fine-grained categories are consolidated into four coarse-grained categories.

In line with previous studies (Yang et al. 2022; Hsu et al. 2023; Song, Sebe, and Wang 2023), top-1 accuracy is used as the evaluation metric on unseen target domains.

### Comparison with SOTA

The compared methods include: 1) generic domain generalization methods that are not especially devised for fine-grained images, namely, ARM (Zhang et al. 2020), DANN (Ganin et al. 2016), MLDG (Li et al. 2018), GroupDRO (Sagawa et al. 2019), CORAL (Sun and Saenko 2016), SagNet (Nam et al. 2021), MixStyle (Zhou et al. 2021), Mixup (Yan et al. 2020), RIDG (Chen et al. 2023), SAGM (Wang et al. 2023), MIRO (Cha et al. 2022), SDViT (Sultana et al. 2022), and GMoE (Li et al. 2023); 2) domain adaptation

Method	C $\rightarrow$ P	P $\rightarrow$ C	Avg	Params
<i>ResNet-Based:</i>				
DP-Net* (DA)	65.22	47.81	56.52	103M
PAN* (DA)	67.40	50.92	59.16	103M
ARM	47.98	31.53	39.76	24M
DANN	54.06	37.09	45.57	24M
MLDG	55.40	34.15	44.78	23M
GroupDRO	54.94	35.67	45.31	23M
CORAL	54.70	35.29	45.00	23M
SagNet	56.33	36.71	46.52	24M
MixStyle	52.97	28.44	40.71	23M
Mixup	54.58	34.66	44.62	23M
RIDG	36.41	24.11	30.26	24M
SAGM	57.83	37.16	47.50	23M
MIRO	56.29	41.28	48.79	47M
FSDG	61.77	48.43	55.10	49M
<i>ViT-Based:</i>				
SDViT	58.43	46.48	52.46	23M
GMoE	60.70	48.13	54.42	34M
<i>VMamba-Based:</i>				
VMamba	65.01	61.93	63.47	33M
<b>HSSH (Ours)</b>	<b>67.48</b>	<b>64.57</b>	<b>66.03</b>	35M

Table 1: Performance comparison of the proposed HSSH and existing methods on the Cub-Paintings dataset. \*: Target domain is available when training.

Method	M $\rightarrow$ N	N $\rightarrow$ M	Avg	Params
<i>ResNet-Based:</i>				
DP-Net* (DA)	59.27	62.01	60.64	103M
PAN* (DA)	59.42	62.15	60.79	103M
ARM	47.84	29.62	38.73	24M
DANN	56.01	35.37	45.69	24M
MLDG	56.92	35.95	46.44	23M
GroupDRO	52.14	34.06	43.10	23M
CORAL	54.47	33.90	44.18	23M
SagNet	57.91	32.58	45.24	24M
MixStyle	53.74	27.45	40.60	23M
Mixup	50.91	36.13	43.52	23M
RIDG	51.95	43.88	47.92	24M
SAGM	56.82	39.59	48.20	23M
MIRO	55.82	39.86	47.84	47M
FSDG	57.35	50.68	54.02	49M
<i>ViT-Based:</i>				
SDViT	54.57	48.76	51.67	23M
GMoE	56.72	50.61	53.67	34M
<i>VMamba-Based:</i>				
Vmamba	60.14	73.55	66.85	33M
<b>HSSH (Ours)</b>	<b>62.60</b>	<b>76.69</b>	<b>69.65</b>	35M

Table 2: Performance comparison of the proposed HSSH and existing methods on the remote sensing dataset. \*: Target domain is available when training.

fine-grained methods, namely, PAN (Du et al. 2021) and DP-Net (Wang et al. 2021); 3) domain generalized fine-grained methods FSDG (Yu et al. 2024).

**Results on CUB-Paintings** are presented in Table 1. The

Method	C → I	C → N	I → C	I → N	N → C	N → I	Avg	Params
<i>ResNet-Based:</i>								
DP-Net* (DA)	68.37	82.66	90.20	87.52	91.47	75.13	82.56	103M
PAN* (DA)	69.79	84.19	90.46	88.10	92.51	75.03	83.34	103M
ARM	50.51	71.25	77.38	74.20	84.74	59.82	69.65	24M
DANN	52.75	71.82	80.79	73.59	85.55	61.53	71.01	24M
MLDG	53.55	72.19	80.74	74.83	85.61	61.95	71.48	23M
GroupDRO	52.61	70.78	81.87	74.40	86.26	61.32	71.21	23M
CORAL	54.64	72.93	81.01	74.97	86.10	62.51	72.03	23M
SagNet	53.66	71.75	81.39	74.13	85.66	62.06	71.44	24M
MixStyle	49.95	69.04	74.46	68.34	83.60	57.12	67.09	23M
Mixup	52.36	71.65	82.36	75.17	85.61	62.34	71.58	23M
RIDG	47.15	66.71	82.47	73.63	85.77	60.98	69.45	24M
SAGM	54.04	73.63	82.96	77.01	87.88	63.49	73.17	23M
MIRO	54.39	74.87	82.36	75.34	86.42	62.48	72.64	47M
FSDG	66.70	83.29	89.97	87.33	92.20	74.34	82.31	49M
<i>ViT-Based:</i>								
SDViT	62.21	79.91	86.58	82.87	88.26	69.64	78.25	23M
GMoE	64.86	82.45	89.07	85.76	90.92	73.59	81.11	34M
<i>VMamba-Based:</i>								
VMamaba	77.98	87.05	92.76	92.00	93.86	85.80	88.24	33M
<b>HSSH (Ours)</b>	<b>80.43</b>	<b>90.39</b>	<b>95.25</b>	<b>94.33</b>	<b>96.17</b>	<b>87.57</b>	<b>90.69</b>	35M

Table 3: Performance comparison of the proposed HSSH and existing methods on the Brids-31 dataset. \*: Target domain is available when training.

Backbone	SSH	HMC	C → P	P → C	Avg
✓			65.01	61.93	63.47
✓	✓		66.42	63.29	64.86
✓	✓	✓	<b>67.48</b>	<b>64.57</b>	<b>66.03</b>

Table 4: Ablation study on each component of the proposed HSSH. Classification accuracy presented in percentage (%).

$F^1$	$F^2$	$F^3$	$F^4$	C → P	P → C	Avg
✓				64.78	61.96	63.53
✓	✓			66.02	62.37	64.20
✓	✓	✓		66.81	63.28	65.05
✓	✓	✓	✓	<b>67.48</b>	<b>64.57</b>	<b>66.03</b>

Table 5: Ablation studies on each stage of SSH. Classification accuracy presented in percentage (%).

proposed HSSH outperforms all other approaches, even surpassing domain adaptation (DA) methods that have access to target domain images. Compared to the second-best FGDG method FSDG, HSSH shows an accuracy improvement of 5.71% and 16.14% in the C→P and P→C settings. Additionally, by expanding the style in the state space and hierarchically alignment in hyperbolic space, HSSH further enhances VMamba’s performance to achieve improvements of 2.47% and 2.64% in the two respective settings.

**Results on RS-FGDG** are presented in Table 2. The remote sensing fine-grained classification task is more challenging compared to CUB-Paintings due to the complexity of spatial distributions and varying resolutions, leading to

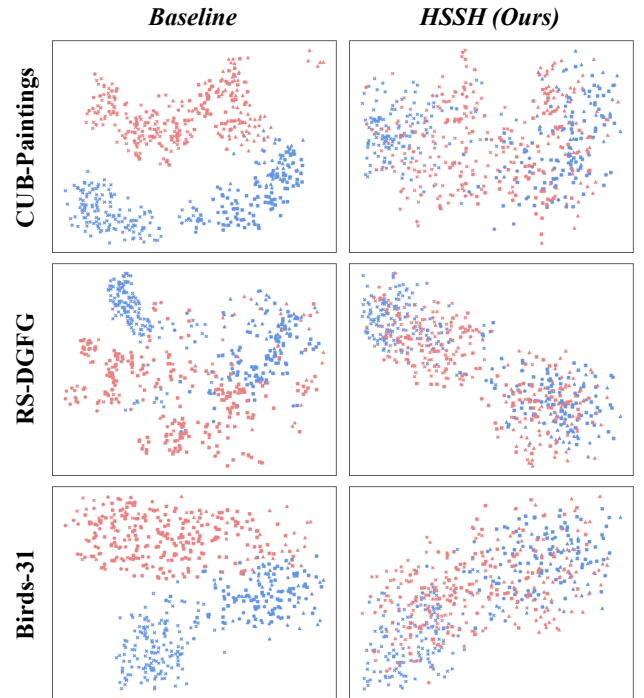


Figure 4: t-SNE feature visualization. Blue: source domain; Red: unseen target domain. Different types of icon refer to different fine-grained categories.

a noticeable limited performance for DG methods. Despite

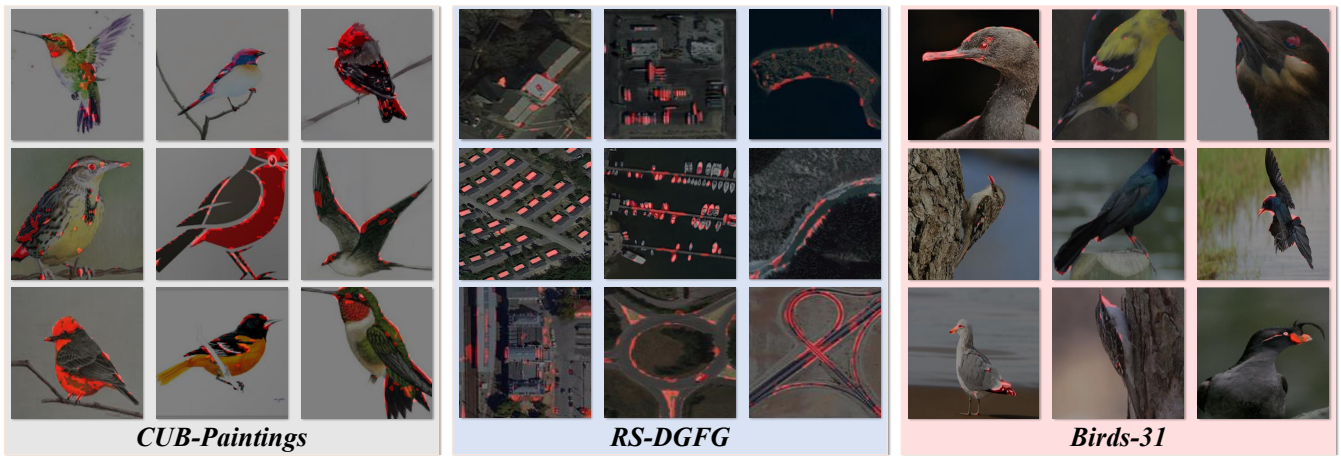


Figure 5: Heatmap visualization of the proposed HSSH on unseen target domains by GradCAM.

this, VMamba achieves the second-best performance in DG method, outperforming FSDG by 5.25% and 26.01% in the  $M \rightarrow N$  and  $N \rightarrow M$  settings, respectively. However, VMamba does not surpass the DA methods DP-Net and PAN. The proposed HSSH scheme demonstrates superior performance, achieving improvements of 2.46% and 3.14% over VMamba and even outperforming DA methods without access to any target domain when training.

**Results on Birds-31** are presented in Table 3. The experiments were conducted across six different settings spanning three domains. The Birds-31 setup involves three natural image datasets with relatively smaller domain gaps, leading to higher performance for both DG and DA methods. The proposed HSSH achieves state-of-the-art results in all settings, outperforming the best DA method PAN, with improvements of 10.64%, 6.20%, 4.79%, 3.36%, 12.54%, and 7.35% for  $C \rightarrow I$ ,  $C \rightarrow N$ ,  $I \rightarrow C$ ,  $I \rightarrow N$ ,  $N \rightarrow C$ , and  $N \rightarrow I$ , respectively. Notably, the average accuracy has been improved by 8.38% compared to the second-best FG DG method FSDG.

### Ablation Study

**On Each Component.** The proposed HSSH consists of the VMamba backbone, state space hallucination module and hyperbolic manifold consistency, which we denote as backbone, SSH and HMC, respectively. For ablation studies, when implementing each component, the corresponding loss function is also implemented. The results under the CUB-Paintings settings are shown in Table 4. Specifically, when SSH is utilized, the classification accuracy improvement is 1.41%, 1.36%, and 1.39% under  $C \rightarrow P$ ,  $P \rightarrow C$ , and the average setting, respectively, compared to using  $\mathcal{L}_{cls}$  alone for FG DG. The use of HMC further improves the classification accuracy by 1.06%, 1.28%, and 1.17%, demonstrating the positive impact of hierarchical hyperbolic embedding on fine-grained representation.

**On Hallucinated Stage.** The hallucinated features  $\tilde{F}^1$ ,  $\tilde{F}^2$ ,  $\tilde{F}^3$ , and  $\tilde{F}^4$  expand the style space at different stages within the state space. An ablation study is conducted on each stage. As shown in Table 5, hallucinating across all scales yields optimal performance on the CUB-Paintings. Notably,

expanding the style space at the final stage  $F^4$  is most critical, resulting in improvements of 0.67%, 1.29%, and 0.98% under the  $C \rightarrow P$ ,  $P \rightarrow C$ , and average setting, respectively.

### Visualization

**On Feature Space.** To inspect if HSSH alleviates the gap between each fine-grained domain over the VMamba baseline, we extract the state embeddings from both the source and unseen target domains. Then, we display them by t-SNE feature visualization, and the results are shown in Fig. 4. For all experiments, blue and red points refer to the samples from the source and unseen target domain, respectively. HSSH allows a more uniform mixture between the source and target domain samples than the VMamba baseline. This indicates its effectiveness to handle the problem of FG DG.

**On Activated Patterns.** Following the visualization of prior FGVC methods, GradCAM is used to visualize the features from the proposed HSSH. The results are displayed in Fig. 5. Overall, the proposed HSSH is able to discern the subtle and tiny fine-grained patterns on the unseen domains from all the three benchmarks, demonstrating its effectiveness.

### Conclusion

In this paper, we propose a novel Hyperbolic State Space Hallucination (HSSH) scheme for fine-grained domain generalization (FGDG). It makes an earlier exploration on not only harnessing state space model (SSM) for FG DG, but also modeling the hyperbolic selective state space. The proposed HSSH consists of two key components, namely, state space hallucination (SSH) and hyperbolic manifold consistency (HMC). SSH enriches the style diversity by firstly extrapolating and then hallucinating the style of the state embeddings. HMC minimizes the hyperbolic distance between the pre- and post- style hallucinate images, which benefits the discernment of fine-grained patterns despite the cross-domain styles. Finally, HMC minimizes the hyperbolic distance, so that the impact from the style variation on fine-grained patterns can be eliminated. Extensive experiments across three FG DG benchmarks demonstrate the state-of-the-art performance of the proposed HSSH.

## References

- Bi, Q.; Yi, J.; Zheng, H.; Ji, W.; Zhan, H.; Huang, Y.; Li, Y.; and Zheng, Y. 2024a. Samba: Severity-aware Recurrent Modeling for Cross-domain Medical Image Grading. In *Annual Conference on Neural Information Processing Systems*.
- Bi, Q.; Yi, J.; Zheng, H.; Zhan, H.; Huang, Y.; Ji, W.; Li, Y.; and Zheng, Y. 2024b. Learning Frequency-Adapted Vision Foundation Model for Domain Generalized Semantic Segmentation. In *Annual Conference on Neural Information Processing Systems*.
- Bi, Q.; You, S.; and Gevers, T. 2024a. Learning content-enhanced mask transformer for domain generalized urban-scene segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 819–827.
- Bi, Q.; You, S.; and Gevers, T. 2024b. Learning generalized segmentation for foggy-scenes by bi-directional wavelet guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 801–809.
- Cha, J.; Lee, K.; Park, S.; and Chun, S. 2022. Domain generalization by mutual-information regularization with pre-trained models. In *European conference on computer vision*, 440–457.
- Chen, L.; Lu, X.; Zhang, J.; Chu, X.; and Chen, C. 2021. Hinet: Half instance normalization network for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 182–192.
- Chen, L.; Zhang, Y.; Song, Y.; Van Den Hengel, A.; and Liu, L. 2023. Domain generalization via rationale invariance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1751–1760.
- Cheng, G.; Han, J.; and Lu, X. 2017. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10): 1865–1883.
- Desai, K.; Nickel, M.; Rajpurohit, T.; Johnson, J.; and Vedantam, S. R. 2023. Hyperbolic Image-text Representations. In *International Conference on Machine Learning*, 7694–7731.
- Du, R.; Xie, J.; Ma, Z.; Chang, D.; Song, Y.-Z.; and Guo, J. 2021. Progressive learning of category-consistent multi-granularity features for fine-grained visual classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ganea, O.; Bécigneul, G.; and Hofmann, T. 2018. Hyperbolic neural networks. *Advances in neural information processing systems*, 31.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; March, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(59): 1–35.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Re, C. 2022. Efficiently Modeling Long Sequences with Structured State Spaces. In *International Conference on Learning Representations*.
- Gu, A.; Johnson, I.; Goel, K.; Saab, K.; Dao, T.; Rudra, A.; and Ré, C. 2021. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *Advances in Neural Information Processing Systems*, 34: 572–585.
- Hsu, Y.-C.; Hong, C.-Y.; Lee, M.-S.; Geiger, D.; and Liu, T.-L. 2023. ABC-Norm Regularization for Fine-Grained and Long-Tailed Image Classification. *IEEE Transactions on Image Processing*, 32: 3885–3896.
- Hu, C.; Zhang, K.-Y.; Yao, T.; Ding, S.; and Ma, L. 2024. Rethinking generalizable face anti-spoofing via hierarchical prototype-guided distribution refinement in hyperbolic space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1032–1041.
- Hu, Y.; Jin, X.; Zhang, Y.; Hong, H.; Zhang, J.; He, Y.; and Xue, H. 2021. RAMS-Trans: Recurrent Attention Multi-scale Transformer for Fine-grained Image Recognition. In *Proceedings of the 29th ACM International Conference on Multimedia*, 4239–4248.
- Huang, X.; and Belongie, S. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, 1501–1510.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Trans. ASME, D*, 82: 35–44.
- Kong, F.; Chen, Y.; Cai, J.; and Modolo, D. 2024. Hyperbolic learning with synthetic captions for open-world detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16762–16771.
- Lee, S.; Seong, H.; Lee, S.; and Kim, E. 2022. WildNet: Learning domain generalized semantic segmentation from the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9936–9946.
- Li, B.; Shen, Y.; Yang, J.; Wang, Y.; Ren, J.; Che, T.; Zhang, J.; and Liu, Z. 2023. Sparse mixture-of-experts are domain generalizable learners. In *International Conference on Learning Representations*.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. 2018. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024. VMamba: Visual state space model. *arXiv preprint arXiv:2401.10166*.
- Long, Y.; Xia, G.-S.; Li, S.; Yang, W.; Yang, M. Y.; Zhu, X. X.; Zhang, L.; and Li, D. 2021. On creating benchmark dataset for aerial image interpretation: Reviews, guidances, and million-aid. *IEEE Journal of selected topics in applied earth observations and remote sensing*, 14: 4205–4230.
- Lu, C.; Schroecker, Y.; Gu, A.; Parisotto, E.; Foerster, J.; Singh, S.; and Behbahani, F. 2024. Structured state space models for in-context reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Mettes, P.; Ghadimi Atigh, M.; Keller-Ressel, M.; Gu, J.; and Yeung, S. 2024. Hyperbolic deep learning in computer vision: A survey. *International Journal of Computer Vision*, 1–25.

- Nam, H.; Lee, H.; Park, J.; Yoon, W.; and Yoo, D. 2021. Reducing domain gap by reducing style bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8690–8699.
- Sagawa, S.; Koh, P. W.; Hashimoto, T. B.; and Liang, P. 2019. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*.
- Song, Y.; Sebe, N.; and Wang, W. 2023. On the Eigenvalues of Global Covariance Pooling for Fine-Grained Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3554–3566.
- Sultana, M.; Naseer, M.; Khan, M. H.; Khan, S.; and Khan, F. S. 2022. Self-distilled vision transformer for domain generalization. In *Proceedings of the Asian Conference on Computer Vision*, 3068–3085.
- Sun, B.; and Saenko, K. 2016. Deep coral: Correlation alignment for deep domain adaptation. In *Europe Conference on Computer Vision Workshop*, 443–450.
- Van Horn, G.; Branson, S.; Farrell, R.; Haber, S.; Barry, J.; Ipeirotis, P.; Perona, P.; and Belongie, S. 2015. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 595–604.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8769–8778.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset.
- Wang, P.; Zhang, Z.; Lei, Z.; and Zhang, L. 2023. Sharpness-aware gradient matching for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3769–3778.
- Wang, S.; Chen, X.; Wang, Y.; Long, M.; and Wang, J. 2020. Progressive adversarial networks for fine-grained domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9213–9222.
- Wang, S.; Li, H.; Wang, Z.; and Ouyang, W. 2021. Dynamic position-aware network for fine-grained image recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2791–2799.
- Weber, M.; Zaheer, M.; Rawat, A. S.; Menon, A. K.; and Kumar, S. 2020. Robust large-margin learning in hyperbolic space. *Advances in Neural Information Processing Systems*, 33: 17863–17873.
- Xu, S.-L.; Sun, Y.; Zhang, F.; Xu, A.; Wei, X.-S.; and Yang, Y. 2023. Hyperbolic space with hierarchical margin boosts fine-grained learning from coarse labels. *Advances in Neural Information Processing Systems*, 36: 71263–71274.
- Yan, S.; Song, H.; Li, N.; Zou, L.; and Ren, L. 2020. Improve unsupervised domain adaptation with mixup training. *ArXiv:2001.00677*.
- Yang, M.; Zhou, M.; Ying, R.; Chen, Y.; and King, I. 2023. Hyperbolic representation learning: Revisiting and advancing. In *International Conference on Machine Learning*, 39639–39659. PMLR.
- Yang, X.; Wang, Y.; Chen, K.; Xu, Y.; and Tian, Y. 2022. Fine-Grained Object Classification via Self-Supervised Pose Alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7399–7408.
- Yi, J.; Bi, Q.; Zheng, H.; Zhan, H.; Ji, W.; Huang, Y.; Li, S.; Li, Y.; Zheng, Y.; and Huang, F. 2024a. Hallucinated style distillation for single domain generalization in medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 438–448.
- Yi, J.; Bi, Q.; Zheng, H.; Zhan, H.; Ji, W.; Huang, Y.; Li, Y.; and Zheng, Y. 2024b. Learning Spectral-decomposed Tokens for Domain Generalized Semantic Segmentation. In *ACM Multimedia 2024*.
- Yu, T.; and De Sa, C. 2019. Numerically accurate hyperbolic embeddings using tiling-based models. In *NeurIPS*.
- Yu, W.; Chen, D.; Wang, Q.; and Hu, Q. 2024. Fine-Grained Domain Generalization with Feature Structuralization. *arXiv preprint arXiv:2406.09166*.
- Zhang, M.; Marklund, H.; Gupta, A.; Levine, S.; and Finn, C. 2020. Adaptive risk minimization: A meta-learning approach for tackling group shift. *arXiv preprint arXiv:2007.02931*, 8(9).
- Zhang, Y.; Li, M.; Li, R.; Jia, K.; and Zhang, L. 2022. Exact feature distribution matching for arbitrary style transfer and domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8035–8045.
- Zhao, Y.; Li, J.; Chen, X.; and Tian, Y. 2021. Part-Guided Relational Transformers for Fine-Grained Visual Recognition. *IEEE Transactions on Image Processing*, 30: 9470–9481.
- Zhao, Y.; Zhong, Z.; Zhao, N.; Sebe, N.; and Lee, G. H. 2022. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In *European Conference on Computer Vision*, 535–552.
- Zhao, Y.; Zhong, Z.; Zhao, N.; Sebe, N.; and Lee, G. H. 2024. Style-hallucinated dual consistency learning: A unified framework for visual domain generalization. *International Journal of Computer Vision*, 132(3): 837–853.
- Zhong, Z.; Zhao, Y.; Lee, G. H.; and Sebe, N. 2022. Adversarial style augmentation for domain generalized urban-scene segmentation. *Advances in Neural Information Processing Systems*, 35: 338–350.
- Zhou, K.; Yang, Y.; Qiao, Y.; and Xiang, T. 2021. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision Mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.