

MBZUAI

Digital.Commons@MBZUAI

Computer Vision Faculty Publications

Scholarly Works

7-12-2021

Structured latent embeddings for recognizing unseen classes in unseen domains

Shivam Chandhok

Mohamed bin Zayed University of Artificial Intelligence

Sanath Narayan

Inception Institute of Artificial Intelligence

Hisham Cholakkal

Mohamed bin Zayed University of Artificial Intelligence

Rao Muhammad Anwer

Mohamed bin Zayed University of Artificial Intelligence

Vineeth N. Balasubramanian

Indian Institute of Technology

See next page for additional authors
See this and additional works at: <https://dclibrary.mbzuai.ac.ae/cvfp>



Part of the [Computer Sciences Commons](#)

Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC BY 4.0

Uploaded 25 March 2022

Recommended Citation

S. Chandhok, et.al., "Structured latent embeddings for recognizing unseen classes in unseen domains", 2021, arXiv:2107.05622.

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Computer Vision Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Authors

Shivam Chandhok, Sanath Narayan, Hisham Cholakkal, Rao Muhammad Anwer, Vineeth N. Balasubramanian, Fahad Shahbaz Khan, and Ling Shao

Structured Latent Embeddings for Recognizing Unseen Classes in Unseen Domains

Shivam Chandhok¹
shivam.chandhok@mbzuai.ac.ae

Sanath Narayan²
sanath.narayan@inceptioniai.org

Hisham Cholakkal¹
hisham.cholakkal@mbzuai.ac.ae

Rao Muhammad Anwer¹
rao.anwer@mbzuai.ac.ae

Vineeth N Balasubramanian⁴
vineethnb@iith.ac.in

Fahad Shahbaz Khan¹³
fahad.khan@mbzuai.ac.ae

Ling Shao²
ling.shao@ieee.org

¹ Mohamed Bin Zayed University
of AI, UAE

² Inception Institute of Artificial
Intelligence, UAE

³ Linköping University, Sweden

⁴ Indian Institute of Technology,
Hyderabad, India

Abstract

The need to address the scarcity of task-specific annotated data has resulted in concerted efforts in recent years for specific settings such as zero-shot learning (ZSL) and domain generalization (DG), to separately address the issues of semantic shift and domain shift, respectively. However, real-world applications often do not have constrained settings and necessitate handling unseen classes in unseen domains – a setting called Zero-shot Domain Generalization, which presents the issues of domain and semantic shifts simultaneously. In this work, we propose a novel approach that learns domain-agnostic structured latent embeddings by projecting images from different domains as well as class-specific semantic text-based representations to a common latent space. In particular, our method jointly strives for the following objectives: (i) aligning the multi-modal cues from visual and text-based semantic concepts; (ii) partitioning the common latent space according to the domain-agnostic class-level semantic concepts; and (iii) learning a domain invariance w.r.t. the visual-semantic joint distribution for generalizing to unseen classes in unseen domains. Our experiments on the challenging Domain-Net and DomainNet-LS benchmarks show the superiority of our approach over existing methods, with significant gains on difficult domains like *quickdraw* and *sketch*.

1 Introduction

In various computer vision problems, obtaining labeled data specifically tailored for a new task (be it a new domain or a new class) can be challenging due to one or more of several reasons: high annotation costs, dynamic addition of objects with new semantic content, limited instances of rare objects or long-tailed distributions which frequently occur in real-world scenarios [80]. To address such issues, two popular recent approaches include: (i) zero-shot learning (ZSL): use training data of related object categories from the same domain (e.g., sketches of cats as a training data for recognizing dogs from sketches); and (ii) domain generalization (DG): use training data of a particular object category from related domains (e.g., photos/real images of dogs as a training data for recognizing dogs from sketches). More recently, there has been increasing interest in a combination of these approaches to handle unseen classes in unseen domains, viz. zero-shot domain generalization (which we call ZSLDG), where one leverages training data of a related object category from a related domain (e.g., photos of cats as a training data for recognizing dogs from sketches). DG addresses only domain shift that occurs due to the training and test domains being different. ZSL addresses semantic shift that occurs due to the presence of different object categories during training and testing. In contrast, ZSLDG more closely aligns with the challenges faced in real-world applications, but needs to simultaneously address domain and semantic shift issues [23, 24].

In this work, we investigate this challenging problem of ZSLDG. In particular, we propose a unified solution that jointly tackles both domain and semantic shifts by relating the general visual cues of a class to semantic concepts that are invariant across domains. E.g., semantic cues such as *<long neck, long legs, has spots>* of *giraffe* class are invariant across domains such as *real* images (photos), *sketch* or *clipart* (see Fig. 1). To this end, we bring common information from visual and semantic spaces into a structured domain-agnostic latent space, partitioned according to class-level semantic concepts. We then impose a domain invariance w.r.t. the visual-semantic joint distribution. Since the semantic space is shared across all classes and is agnostic to the visual domains, imposing such invariance aids in generalizing to unseen domains at test time (instead of overfitting to source domains), while improving the visual-semantic interaction for effective knowledge transfer across seen and unseen classes.

Contributions: We propose a ZSLDG approach comprising of a visual encoder that learns to project multi-domain images from the visual space to a latent space, and a semantic encoder that learns to map text-based category-specific semantic representations to the same latent space. The key contributions of the proposed approach are: (i) For aligning class-specific cues from visual and semantic latent embeddings, we introduce a multimodal alignment loss term; (ii) We propose to partition the latent space w.r.t. class-level semantic concepts across domains by minimizing intra-class variance across different seen domains; (iii) The focus

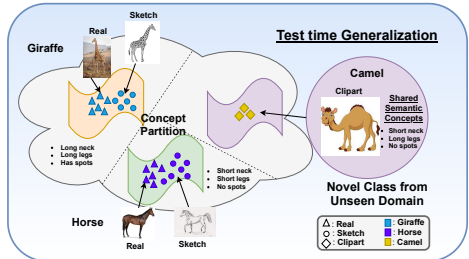


Figure 1: Our latent space is structured according to class-level semantic concepts and is domain-invariant w.r.t. visual-semantic joint distribution. This enables our model to map unseen classes in unseen domains at test-time (*camel* from *clipart*), based on their semantic concepts, to appropriate subspaces within our latent space, thereby aiding generalization.

of our design is introduction of a joint invariance module that seeks to achieve domain invariance w.r.t. the visual-semantic joint distribution, and thereby facilitates generalizing to unseen classes in unseen domains; and (iv) Experiments and ablation studies on the challenging DomainNet and DomainNet-LS benchmarks [31] demonstrate the superiority of our approach over existing methods. Particularly, on the most difficult *quickdraw* domain, our approach achieves a significant gain of 1.6% over the best existing method [23].

2 Related Work

Domain Generalization (DG): Existing methods tackle the problem of domain shift, which occurs when the training and testing data belong to different domains, in different ways. Most previous approaches aim to learn domain-invariance by minimizing the discrepancy between multiple source domains [27, 23, 44] or by employing autoencoders and adversarial losses [12, 19]. A few works [8, 17, 18] introduce specific training policies or optimization procedures such as meta-learning and episodic training to enhance the generalizability of the model to unseen domains. Similarly, [24, 37] employ data augmentation strategies to improve the models robustness to data distribution shifts at test time. However, all these works tackle the DG problem alone, where the label spaces at both train and test time are identical.

Zero-shot Learning (ZSL): Traditional ZSL methods [9, 8, 9, 32, 35] learn to project the visual features onto a semantic embedding space via direct mapping or through a compatibility function. However, such direct mappings are likely to suffer from issues of seen class bias and hubness [9, 14]. In contrast, the work of [36] leverages joint multi-modal learning of visual and textual feature embeddings for the task of ZSL. Recently, generative approaches tackle the problem of seen class bias by generating unseen visual features from respective class embeddings [8, 21, 76, 28, 33, 38, 40, 42]. However, all the aforementioned methods address only ZSL, where the domain remains unchanged during training and testing.

Zero-shot Domain Generalization (ZSLDG): Recently, CuMix [23] introduced the problem of ZSLDG. While [24] defined variations in rotations of the same objects as different domains, such a restricted definition limits its real-world applicability. Differently, CuMix [23] defines domains as different ways of depicting an object, as in *sketch, painting, cartoon, etc.*, which is closer to practical use of such methods. CuMix tackles the issue of domain shift through data augmentation by mixing and interpolating source domains, and handles semantic shifts by learning to project visual features to the semantic space. This work also established a benchmark dataset, DomainNet, for this setting with an evaluation protocol, which we follow in this work for fair comparison. However, relying on mixing source domains has a drawback – the resulting model could overfit to the source domains and their interpolations, thereby reducing generalizability to unseen domains [20]. Furthermore, directly mapping the visual space to the semantic space, as in [23], can lead to hubness issues (mapped points cluster as a hub due to low variance) [9, 14], thereby reducing class-discriminative capability. In contrast, our approach jointly handles the issues of domain and semantic shifts by learning a domain-agnostic latent space that is partitioned based on class-level (domain-invariant) semantic concepts, onto which the visual and semantic features are projected. Since domain invariance is enforced w.r.t. the visual-semantic joint distribution, it is less likely to overfit to seen domains (a common problem when domain-invariance is enforced w.r.t. marginal distribution of images [20, 22]). In addition, our approach enables better interaction between visual and semantic spaces in a new latent space, thereby supporting model generalization to unseen classes in unseen domains.

3 Proposed Method

Problem Setting: The goal in zero-shot domain generalization (ZSLDG) is to recognize unseen categories in unseen domains. Let $Q^{Tr} = \{(\mathbf{x}, y, \mathbf{a}_y, d) | \mathbf{x} \in \mathcal{X}, y \in \mathcal{Y}^s, \mathbf{a}_y \in \mathcal{A}, d \in \mathcal{D}^s\}$ denote the training set, where \mathbf{x} is a seen class image in the visual space (\mathcal{X}) with corresponding label y from a set of seen class labels \mathcal{Y}^s . Here, \mathbf{a}_y denotes the class-specific semantic representation that encodes the inter-class relationships, while d is the domain label from a set of seen domains \mathcal{D}^s . Note that the semantic representations are typically obtained from unsupervised text-based WordNet models (e.g., *word2vec* [25]). Similarly, $Q^{Ts} = \{(\mathbf{x}, y, \mathbf{a}_y, d) | \mathbf{x} \in \mathcal{X}, y \in \mathcal{Y}^u, \mathbf{a}_y \in \mathcal{A}, d \in \mathcal{D}^u\}$ is the test set, where \mathcal{Y}^u is the set of labels for unseen classes and \mathcal{D}^u represents the set of unseen domains. In the standard zero-shot setting, images at training and testing belong to disjoint classes but share the same domain space, i.e., $\mathcal{Y}^s \cap \mathcal{Y}^u \equiv \emptyset$ and $\mathcal{D}^s \equiv \mathcal{D}^u$. On the other hand, in the standard DG setting, images at training and testing belong to same categories in disjoint domain spaces, i.e., $\mathcal{Y}^s \equiv \mathcal{Y}^u$ and $\mathcal{D}^s \cap \mathcal{D}^u \equiv \emptyset$. In this work, our goal is to address the more challenging ZSLDG setting for recognizing unseen classes in unseen domains without having seen these novel classes and domains during training, i.e., $\mathcal{Y}^s \cap \mathcal{Y}^u \equiv \emptyset$ and $\mathcal{D}^s \cap \mathcal{D}^u \equiv \emptyset$.

Overall Framework: The overall architecture of our proposed approach is shown in Fig. 2. The proposed framework comprises a visual encoder f , semantic encoder g , semantic projection classifier h along with discriminators D_1 and D_2 . In ZSLDG, the conditional distribution $p(y|\mathbf{x})$ changes since \mathbf{x} comes from different domains, i.e., $p_{\mathcal{X}}(\mathbf{x}|d_i) \neq p_{\mathcal{X}}(\mathbf{x}|d_j), \forall i \neq j$. Our approach mitigates this issue by learning a domain-invariant semantic manifold \mathcal{Z} which is partitioned according to class-level semantic concepts (described in Sec. 3.1 and Sec. 3.2), such that $p(y|\mathbf{z})$ is stable and does not change across domains (where $\mathbf{z} = f(\mathbf{x})$). Furthermore, in order to ensure generalization to unseen classes in unseen domains at test time, our joint invariance module achieves domain invariance w.r.t. the visual-semantic joint by employing $\mathcal{L}_{joint-inv}$ (described in Sec. 3.3). This facilitates improved knowledge transfer between class-specific (domain-invariant) visual cues and semantic representations in latent space \mathcal{Z} , thereby enhancing generalization to unseen classes in unseen domains at test-time.

3.1 Multimodal Alignment

The multimodal alignment module, learns to project both the visual and semantic representations to a common latent embedding space \mathcal{Z} . Let $f(\mathbf{x}) : \mathcal{X} \rightarrow \mathcal{Z}$ denote a feature extractor, which maps an image \mathbf{x} in the visual space \mathcal{X} to a vector \mathbf{z}_v in the latent embedding space \mathcal{Z} . Furthermore, let the function g learn a mapping from semantic space to the latent embedding space, i.e., $g(\mathbf{n}, \mathbf{a}_y) : \mathcal{N} \times \mathcal{A} \rightarrow \mathcal{Z}$ by taking a random Gaussian noise vector \mathbf{n} concatenated with the semantic representation \mathbf{a}_y as input and mapping it to a vector \mathbf{z}_a in \mathcal{Z} . Let $D_1 : \mathcal{Z} \times \mathcal{A} \rightarrow \mathbb{R}$ denote a conditional discriminator (conditioned on the semantic embedding \mathbf{a}_y). Then, the multimodal adversarial alignment of the visual and semantic embedding spaces is achieved by employing a Wasserstein GAN [9], as given by

$$\mathcal{L}_{D_1} = \mathbb{E}[D_1(\mathbf{z}_v, \mathbf{a}_y)] - \mathbb{E}[D_1(\mathbf{z}_a, \mathbf{a}_y)] - \lambda \mathbb{E}[(\|\nabla_{\tilde{\mathbf{z}}} D_1(\tilde{\mathbf{z}}, \mathbf{a}_y)\|_2 - 1)^2], \quad (1)$$

where $\mathbf{z}_v = f(\mathbf{x})$ and $\mathbf{z}_a = g(\mathbf{n}, \mathbf{a}_y)$ are the latent embeddings from the visual and semantic spaces, respectively. Here, λ is a weighting coefficient, while $\tilde{\mathbf{z}} = \eta \mathbf{z}_v + (1 - \eta) \mathbf{z}_a$ with $\eta \sim U(0, 1)$ represents a convex combination of \mathbf{z}_v and \mathbf{z}_a . Eq. 1 is equivalent to minimizing the (forward) Kullback-Leibler (KL) divergence between the visual and semantic latent embeddings, i.e., $KL[(\mathbf{z}_v, \mathbf{a}_y) || (\mathbf{z}_a, \mathbf{a}_y)]$. Furthermore, to enhance the discriminability of learned

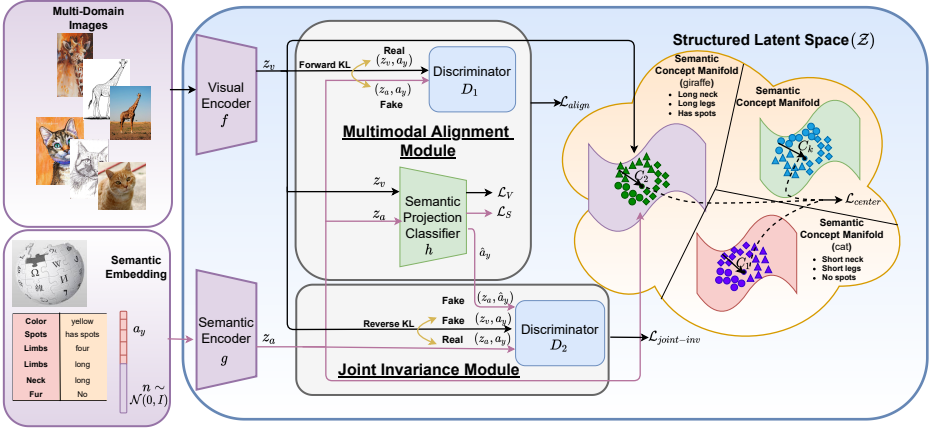


Figure 2: Overall architecture of our approach. The proposed approach comprises a visual encoder f and a semantic encoder g . The multimodal alignment module (Sec. 3.1) aligns the class-specific cues from the visual and semantic latent embeddings (\mathbf{z}_v and \mathbf{z}_a) in \mathcal{Z} by employing an alignment loss term \mathcal{L}_{align} . The loss term \mathcal{L}_{center} ensures a domain-agnostic class-level partitioning (Sec. 3.2) of \mathcal{Z} . Furthermore, the joint invariance module strives to achieve domain invariance (Sec. 3.3) w.r.t. the visual-semantic joint distribution by employing $\mathcal{L}_{joint-inv}$, thereby enabling us to generalize to unseen classes in unseen domains.

latent embeddings, we employ a compatibility based classifier using a semantic projection function $h: \mathcal{Z} \rightarrow \mathcal{A}$ for constraining the latent embeddings (\mathbf{z}_v and \mathbf{z}_a) to map back to their corresponding semantic representations \mathbf{a}_y , given by,

$$\mathcal{L}_V(\mathbf{z}_v, \mathbf{a}_y) = -\mathbb{E}(\log \frac{\exp(\langle h(\mathbf{z}_v), \mathbf{a}_y \rangle)}{\sum_{y \in Y^s} \exp(\langle h(\mathbf{z}_v), \mathbf{a}_y \rangle)}), \quad \mathcal{L}_S(\mathbf{z}_a, \mathbf{a}_y) = -\mathbb{E}(\log \frac{\exp(\langle h(\mathbf{z}_a), \mathbf{a}_y \rangle)}{\sum_{y \in Y^s} \exp(\langle h(\mathbf{z}_a), \mathbf{a}_y \rangle)}). \quad (2)$$

Here, $\langle \cdot, \cdot \rangle$ represents the measure of similarity between its inputs, computed as the dot product between them. Such a cyclic projection, *i.e.*, mapping from visual/semantic space to a latent space and then back to the semantic space minimizes the information loss and enhances the discriminability of the latent embeddings. We employ the multimodal alignment loss term (\mathcal{L}_{align}) to learn the visual and semantic encoders along with the semantic projection classifier, given by

$$\mathcal{L}_{align} = \mathbb{E}[D_1(\mathbf{z}_v, \mathbf{a}_y)] - \mathbb{E}[D_1(\mathbf{z}_a, \mathbf{a}_y)] + \mathcal{L}_V(\mathbf{z}_v, \mathbf{a}_y) + \mathcal{L}_S(\mathbf{z}_a, \mathbf{a}_y). \quad (3)$$

3.2 Structured Partitioning

While the multimodal alignment aligns the visual and corresponding semantic embeddings in the latent space, it does not learn a domain-agnostic latent space, which is partitioned according to the semantic concepts that relate to the different classes. In order to achieve a structured and domain-invariant latent space, we propose to cluster the latent embeddings based on class-level (domain-invariant) semantic concepts across different domains. The latent space is then conceptually structured, since the visual latent embeddings \mathbf{z}_v and semantic latent embeddings \mathbf{z}_a of a class are clustered together. To this end, we adopt the center loss [69] in a multimodal setting. Formally, we first randomly initialise S centers, *i.e.*, $\{\mathbf{c}_j | j = 1, \dots, S\}$ for each of the seen classes in the training set and compute the loss, \mathcal{L}_{center} due to each class y present in a mini-batch. Then, for every class y that is present in a

mini-batch, the center update $\Delta \mathbf{c}_y$ is computed for incrementing the corresponding center \mathbf{c}_y . The loss \mathcal{L}_{center} and update $\Delta \mathbf{c}_y$ are given by:

$$\mathcal{L}_{center} = \delta[\mathbb{E}(\|\mathbf{z}_v - \mathbf{c}_y\|_2^2) + \mathbb{E}(\|\mathbf{z}_a - \mathbf{c}_y\|_2^2)]; \quad \Delta \mathbf{c}_y = \mathbb{E}[\mathbf{c}_y - \mathbf{z}_v] + \mathbb{E}[\mathbf{c}_y - \mathbf{z}_a]. \quad (4)$$

Here, \mathbf{c}_y denotes the center of class label y in the latent space, while \mathbf{z}_v and \mathbf{z}_a correspond to the visual and semantic embeddings of class y , and δ is weighing factor for center loss. Consequently, the intra-class and inter-domain variances for each class get minimized, resulting in a structured and domain-agnostic latent space. Furthermore, since both the visual and semantic latent representations of a class are clustered together, the latent space is partitioned based on class-level semantic concepts.

In order to validate our hypotheses that a domain-agnostic structured latent space helps to stabilize $p(y|\mathbf{z})$ and generalize to new domains, we conduct an experiment as a proof of concept. Fig. 3 presents a comparison for the standard domain generalization (DG) setting on the PACS dataset [16] using ResNet-18 backbone. We see that structuring the latent space (blue bars) provides performance gains on all domains and enhances the average gain, compared to employing multimodal alignment alone (orange bars). The highest gain is achieved for the most difficult *sketch* domain that has a large domain shift from the source domains (*photo*, *art*, *cartoon*), demonstrating the advantage of our domain-agnostic partitioning.




Figure 3: Impact of our structured partitioning for the DG task on PACS [16]. Compared to multimodal alignment alone (orange bars), additionally partitioning the latent space according to the semantic concepts along with multimodal alignment provides notable performance gains (blue bars), especially on the most difficult unseen domain, *i.e.*, *sketch*.

3.3 Joint Invariance Module

As discussed above, the multimodal alignment and conceptual partitioning result in a structured and domain-agnostic latent embedding space that disentangles semantic and domain-specific information. Such a disentanglement of semantic and domain-specific information is sufficient for standard domain-generalization setting where images during training and testing come from same categories. However in our ZSLDG setting, the disentanglement may not hold for unseen semantic categories during testing, as previously found in [17]. In order to address this issue and enable generalization to unseen classes in unseen domains, we propose to learn the domain-invariance w.r.t. the joint distribution of visual and semantic representations of a class. Formally, any given image \mathbf{x} comprises of a class-specific content \mathbf{C} and a domain-specific transformation $T(\cdot)$ which depicts the class in that particular domain d . Thus, each image $\mathbf{x} \in \mathcal{X}$ belonging to domain d_i can be represented as $\mathbf{x} = T_i(\mathbf{C})$. In order to enable generalization to unseen class in unseen domains, we propose to match the visual-semantic joint distribution $p(T(\mathbf{C}), \mathbf{a}_y)$ under different transformations $T_i(\cdot)$ (or domains). Since the semantic space is shared between seen and unseen classes, learning domain invariance w.r.t. the joint distribution of visual and semantic representations of a class, *i.e.*, $p(f(T(\mathbf{C})), \mathbf{a}_y)$ or $p(f(\mathbf{x}), \mathbf{a}_y)$ enables us to enhance generalization.

Specifically, we aim to match the visual-semantic joint distribution from the visual encoder ($\mathbf{z}_v, \mathbf{a}_y$), semantic encoder ($\mathbf{z}_a, \mathbf{a}_y$) and projection classifier ($\mathbf{z}_a, \hat{\mathbf{a}}_y$). To this end, we employ a triple adversarial loss, to stabilize the visual-semantic joint distribution across different domains. This also enhances visual-semantic interaction for learning class-specific

discriminative features in the visual and semantic embedding spaces. This is achieved by employing a discriminator $D_2 : \mathcal{Z} \times \mathcal{A} \rightarrow \mathbb{R}$ and optimizing:

$$\begin{aligned} \mathcal{L}_{D_2} = & \mathbb{E}[D_2(\mathbf{z}_a, \mathbf{a}_y)] - \alpha \mathbb{E}[D_2(\mathbf{z}_a, \hat{\mathbf{a}}_y)] - \beta \mathbb{E}[D_2(\mathbf{z}_v, \mathbf{a}_y)] \\ & - \lambda \mathbb{E}[(\|\nabla_{\tilde{\mathbf{z}}} D_2(\tilde{\mathbf{z}}, \tilde{\mathbf{a}}_y), \nabla_{\tilde{\mathbf{a}}} D_2(\tilde{\mathbf{z}}, \tilde{\mathbf{a}}_y)\|_2 - 1)^2] \end{aligned} \quad (5)$$

Here, $\hat{\mathbf{a}}_y = h(\mathbf{z}_a)$ is output from projection classifier h , which represents the projection of the latent embedding \mathbf{z}_a onto the semantic space \mathcal{A} . Also, $\tilde{\mathbf{z}} = \eta \mathbf{z}_a + (1 - \eta)(\alpha \mathbf{z}_a + \beta \mathbf{z}_v)$ and $\tilde{\mathbf{a}}_y = \eta \mathbf{a}_y + (1 - \eta)(\alpha \hat{\mathbf{a}}_y + \beta \mathbf{a}_y)$ with $\beta = 1 - \alpha$ and $\eta \sim U(0, 1)$. Additionally, λ is a weighting coefficient. Note that D_2 is different from the vanilla discriminator D_1 and has a triple adversarial formulation [15]. Firstly, by incorporating the projection classifier output $\hat{\mathbf{a}}_y$, it enables to jointly train the visual encoder f , semantic encoder g and projection classifier h while imposing domain-invariance. In addition, we design Eq. 5 to treat $(\mathbf{z}_a, \mathbf{a}_y)$ as real samples and $(\mathbf{z}_v, \mathbf{a}_y)$, $(\mathbf{z}_a, \hat{\mathbf{a}}_y)$ as fake samples. This acts as a minimizer of the reverse KL divergence *i.e.*, $KL[(\mathbf{z}_a, \mathbf{a}_y) \| (\mathbf{z}_v, \mathbf{a}_y)]$ [29] (in contrast to D_1 that minimizes forward KL as described in Sec. 3.1) between the visual and semantic spaces. We find that this leads to better generalization by alleviating the mode collapse issue, and thus enables our model to capture multiple modes of the data distribution [29]. Next, the semantic projector classifier h is updated to minimize:

$$\mathcal{L}_{cls} = -\alpha \mathbb{E}[p_h(y|\mathbf{z}_a) D_2(\mathbf{z}_a, \hat{\mathbf{a}}_y)] + \gamma [\mathcal{L}_V(\mathbf{z}_v, \mathbf{a}_y) + \mathcal{L}_S(\mathbf{z}_a, \mathbf{a}_y)], \quad (6)$$

where $p_h(y|\mathbf{z}_a)$ is the probability distribution after taking softmax of semantic projection classifier h , output logits. Weighting the D_2 output with the class probabilities helps in achieving stable training [15]. Finally, we update the visual and semantic encoders (f and g) to minimize discrepancy between the embeddings (\mathbf{z}_a and \mathbf{z}_v) in the latent space, given by:

$$\mathcal{L}_{gen} = \mathbb{E}[D_2(\mathbf{z}_a, \mathbf{a}_y)] - \beta \mathbb{E}[D_2(\mathbf{z}_v, \mathbf{a}_y)]. \quad (7)$$

Then, the joint invariance loss term $\mathcal{L}_{joint-inv}$ is defined as $\mathcal{L}_{joint-inv} = \mathcal{L}_{cls} + \mathcal{L}_{gen}$. Consequently, the adversarial loss terms in Eq. 5 and $\mathcal{L}_{joint-inv}$ together enable us to jointly train f, g, h and learn a domain-invariant space, which can generalize to unseen domains and classes at test time, by capturing class-specific discriminative visual-semantic relationships across domains.

3.4 Training and Inference

Training: In a single training iteration, we first update the discriminators D_1 and D_2 to maximize the losses in Eq. 1 and 5. We update the discriminators 5 times for every update of the rest of the functions (f, g, h), as in WGAN [13]. Following this, the parameters $\theta_f, \theta_g, \theta_h, \theta_c$ corresponding to f, g, h and class centers, respectively, are updated to minimize:

$$\mathcal{L}_{total} = \mathcal{L}_{align} + \mathcal{L}_{center} + \mathcal{L}_{joint-inv}. \quad (8)$$

Inference: A test image \mathbf{x}_t from a unseen domain and class (in \mathcal{D}^u and \mathcal{Y}^u) is projected by encoder f to obtain the corresponding latent embedding $\mathbf{z}_t = f(\mathbf{x}_t)$. The semantic projection classifier h computes pairwise similarities between \mathbf{z}_t and the unseen class embeddings \mathbf{a}_y , where $y \in \mathcal{Y}^u$. These similarity scores are converted to class probabilities to obtain the final prediction \hat{y} , given by $\hat{y} = \arg \max_{y \in \mathcal{Y}^u} P(y|\mathbf{x}_t; \Phi)$.

Table 1: State-of-the-art comparison for the task of ZSLDG on the DomainNet benchmark using ResNet-50 backbone [23]. For a fair comparison, all reported results employ the same backbone, protocol and splits, as described in [23]. Best results are in bold.

DG	Method	AVG	Target Domain				
			<i>painting</i>	<i>infograph</i>	<i>quickdraw</i>	<i>sketch</i>	<i>clipart</i>
-	DEVICE [40]	14.4	17.6	11.7	6.1	16.7	20.1
	ALE [8]	16.2	20.2	12.7	6.8	18.5	22.7
	SPNet [40]	19.4	23.8	16.9	8.2	21.8	26.0
DANN [40]	DEVICE [40]	13.9	16.4	10.4	7.1	15.1	20.5
	ALE [8]	15.7	19.7	12.5	7.4	17.9	21.2
	SPNet [40]	19.1	24.1	15.8	8.4	21.3	25.9
EpiFCR [40]	DEVICE [40]	15.9	19.3	13.9	7.3	17.2	21.6
	ALE [8]	17.5	21.4	14.1	7.8	20.9	23.2
	SPNet [40]	20.0	24.6	16.7	9.2	23.2	26.4
CuMix(Mixup-img-only)		19.2	24.4	16.3	8.7	21.7	25.2
CuMix(Mixup-two-level)		19.9	25.3	17	8.8	21.9	26.6
CuMix [40]		20.7	25.5	17.8	9.9	22.6	27.6
Ours		21.9	26.6	18.4	11.5	25.0	27.8

Table 2: Results on DomainNet-LS with only *real* and *painting* as source domains and ResNet-50 backbone, following protocol described in [23]. Best results in bold.

Model	AVG	<i>quickdraw</i>	<i>sketch</i>	<i>infograph</i>	<i>clipart</i>
SPNet	14.4	4.8	17.3	14.1	21.5
Epi-FCR+SPNet	15.4	5.6	18.7	14.9	22.5
CuMix (MixUp-img-only):	14.3	4.8	17.3	14.0	21.2
CuMix (MixUp-two-level):	15.8	4.9	19.1	16.5	22.7
CuMix (reverse):	15.4	4.8	18.2	15.8	22.9
CuMix:	16.5	5.5	19.7	17.1	23.7
Ours	16.9	7.2	20.5	16	24

Table 3: Ablation study for different components of our framework on DomainNet dataset for ZSLDG setting. Best results are in bold.

Model	AVG	<i>painting</i>	<i>infograph</i>	<i>quickdraw</i>	<i>sketch</i>	<i>clipart</i>
M1: \mathcal{L}_{align}	18.5	22.6	16.2	9.6	20.8	23.7
M2: M1 + \mathcal{L}_{center}	20.5	25.4	16.9	9.8	24.0	26.4
M3: M2 + $\mathcal{L}_{joint-img}$	21.9	26.6	18.4	11.5	25.0	27.8

4 Experiments

Datasets: We evaluate our method on the DomainNet and DomainNet-LS benchmarks for the task of ZSLDG, as in [23]. **DomainNet [30]:** It is a large-scale dataset and is currently the only benchmark dataset for the ZSLDG setting [23]. It consists of nearly 0.6 million images from 345 categories in 6 domains: *painting*, *clipart*, *sketch*, *infograph*, *quickdraw* and *real*. For the task of ZSLDG, we follow the same training/validation/testing splits along with the training and evaluation protocol described in [23]. In particular, 45 out of 345 are fixed as unseen classes and training is performed using only the remaining seen class images. Among the 6 domains in DomainNet, the seen class images from 5 domains are provided during training, and the model is evaluated on the 45 unseen classes in the held-out (unseen) domain. We repeat experiments with each of the domains as the unseen domain. Following [23], the *real* domain is never held out since a ResNet-50 backbone, pre-trained on ImageNet [8], is employed. Average per-class accuracy is used as the performance metric for evaluation on the held-out domain. Similarly, we use the *word2vec* [25] representations as the semantic information for inter-relating seen and unseen classes, as in [23].

DomainNet-LS: This benchmark is a more challenging setting, where the source domains during training are limited to *real* and *painting* only, whereas testing is conducted on the remaining four unseen domains. Since only two source domains are used in training, it is more challenging to learn domain-invariance and generalize at test-time.

4.1 Results: Comparison with State-of-the-art

Results on DomainNet: Tab. 1 shows the comparison of our proposed framework with state-of-the-art methods and all baselines, as established in [23], on the ZSLDG task. We first report the performance of standalone ZSL approaches such as DEVICE [40], ALE [8] and SPNet [40] on the ZSLDG task, followed by the performance achieved by coupling

these ZSL approaches with standard DG approaches like DANN [10] and EpiFCR [18]. It is worth noting that coupling the standalone ZSL methods with DANN achieves lower performance than the ZSL method alone in the case of ZSLDG, since standard domain alignment methods have been shown to be ineffective on the DomainNet dataset, leading to negative transfer in some cases [30]. Furthermore, as noted by [23], coupling EpiFCR (a standalone DG method) with the standalone ZSL approaches is not straightforward, since it requires careful adaptation that includes re-structuring of the loss terms. In particular, the approach of EpiFCR+SPNet achieves an average accuracy (AVG) of 20.0 over different target domains. The recently introduced CuMix [23] approach that targets ZSLDG, employs a curriculum-based mixing policy to generate increasingly complex training samples by mixing up multiple seen domains and categories available during training. The current state-of-art CuMix improves ZSLDG performance over EpiFCR+SPNet, achieving an average accuracy of 20.7 across the target domains. Our approach outperforms CuMix with an absolute gain of 1.2% average across domains ($\sim 6\%$ relative increase) and achieves average accuracy of 21.9 across the five target domains, setting a new state of the art. Furthermore, our method achieves consistent gains over CuMix on each of the target domains.

Results on DomainNet-LS: Tab. 2 shows the performance comparison on the DomainNet-LS benchmark. The SPNet [41] (for standard ZSL) achieves an average accuracy of 14.4, while its integration with EpiFCR [18] (a standard DG approach) improves the performance to 15.4. The current state-of-art CuMix [23] approach for ZSLDG, achieves 16.5 as the average accuracy across the unseen domains. Despite the limited information available during training (and higher domain shift at test time), our approach improves over CuMix by achieving an average accuracy of 16.9 (1.6% relative gain), thereby showing better generalization.

4.2 Ablation Study

We perform an ablation study to understand the efficacy of each component in our proposed method for the ZSLDG task. Tab. 3 shows the performance gains achieved (on DomainNet [61]) by integrating one contribution at a time, in our approach, as below:

- The model learned by employing our multimodal alignment loss term \mathcal{L}_{align} alone (detailed in Sec. 3.1) is denoted as M1
- Similarly, M2 denotes the model learned by integrating \mathcal{L}_{align} with our loss term \mathcal{L}_{center} , which achieves a structured latent space (Sec. 3.2).
- M3 denotes our overall framework, which is learned by integrating our joint invariance loss term $\mathcal{L}_{joint-inv}$ (Sec. 3.3) with \mathcal{L}_{align} and \mathcal{L}_{center} .

The M1 model, which performs multimodal adversarial alignment achieves an average accuracy (denoted as AVG in Tab. 3) of 18.5 across the target domains. Learning a structured latent embedding space along with the multimodal alignment enables the M2 model to achieve an average gain of 2.0 over M1 on the target domains. We note that the gains in M2 due to the integration of \mathcal{L}_{center} with \mathcal{L}_{align} are considerably high on the easier target domains (*clipart*, *painting* and *sketch*). This suggests that \mathcal{L}_{center} is able to achieve an improved structuring of the latent embedding space. Our overall framework (M3) obtains the best results by achieving an average accuracy of 21.9 on the five target domains. Since M3 additionally involves learning the domain invariance w.r.t. the visual-semantic joint by employing $\mathcal{L}_{joint-inv}$, it aids in improving ZSLDG performances on harder target domains such as *quickdraw* and *infograph*. These results clearly indicate that along with the multimodal alignment (\mathcal{L}_{align}), structuring the latent space (\mathcal{L}_{center}) and learning the domain invariance w.r.t. the visual-semantic joint ($\mathcal{L}_{joint-inv}$) are important for recognizing unseen classes in

unseen domains.

5 Conclusions

We propose a novel approach to address the challenging problem of recognizing unseen classes in unseen domains (ZSLDG). Our method learns a domain-agnostic structured latent embedding space which is achieved by employing a multimodal alignment loss term that aligns the visual and semantic spaces, a center loss term that separates different classes in the latent space and a joint invariance term that aids in handling new classes from unseen domains. Our experiments and ablation studies on challenging benchmarks (DomainNet, DomainNet-LS) show the superiority of our approach over existing methods. Future directions include leveraging self-supervision to obtain domain-invariant features and tackle dynamic changes in the label space of categories.

References

- [1] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid. Label-embedding for attribute-based classification. In *CVPR*, pages 819–826, 2013.
- [2] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele. Evaluation of output embeddings for fine-grained image classification. *CVPR*, 2015.
- [3] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid. Label embedding for image classification. *TPAMI*, 2016.
- [4] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein gan. *ICML*, 2017.
- [5] Y. Balaji, S. Sankaranarayanan, and R. Chellappa. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*, 2018.
- [6] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009.
- [7] G. Dinu and M. Baroni. Improving zero-shot learning by mitigating the hubness problem. *CoRR*, abs/1412.6568, 2015.
- [8] R. Felix, V. BG Kumar, I. Reid, and G. Carneiro. Multi-modal cycle-consistent generalized zero-shot learning. *ECCV*, 2018.
- [9] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, M. A. Ranzato J. Dean, and T. Mikolov. Devise: A deep visual-semantic embedding model. *NeurIPS*, 2013.
- [10] A. Frome, G.S. Corrado, J. Shlens, S.Bengio, J.Dean, M. Ranzato, and T.Mikolov. Devise: A deep visual-semantic embedding model. In *NeurIPS*, pages 2121–2129, 2013.
- [11] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Lavioletteand, M.Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(1):2096–2030, 2016.
- [12] M. Ghifary, W. Kleijn, M. Zhang, and D. Balduzzi. Domain generalization for object recognition with multi-task autoencoders. (*ICCV*), 2015.

- [13] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. Improved training of wasserstein gans. *NeurIPS*, 2017.
- [14] A. Lazaridou, G. Dinu, and M. Baroni. Hubness and pollution: Delving into crossspace mapping for zero-shot learning. *ACL*, 2015.
- [15] C. Li, K. Xu, J. Zhu, and B. Zhang. Triple generative adversarial nets. In *NeurIPS*, 2017.
- [16] D. Li, Y. Yang, Y. Song, and T. M. Hospedales. Deeper, broader and artier domain generalization. (*ICCV*), pages 5543–5551, 2017.
- [17] D. Li, Y. Yang, Y. Song, and T. M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018.
- [18] D. Li, J. Zhang, Y. Yang, C. Liu, Y. Song, and T. M. Hospedales. Episodic training for domain generalization. (*ICCV*), pages 1446–1455, 2019.
- [19] H. Li, S. J. Pan, S. Wang, and A. Kot. Domain generalization with adversarial feature learning. *CVPR*, pages 5400–5409, 2018.
- [20] H. Li, S. J. Pan, S. Wang, and A. Kot. Domain generalization with adversarial feature learning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018.
- [21] J. Li, M. Jing, K. Lu, Z. Ding, L. Zhu, and Z. Huang. Leveraging the invariant side of generative zero-shot learning. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7394–7403, 2019.
- [22] Y. Li, X. Tian, M. Gong, Y. Liu, T. Liu, K. Zhang, and D. Tao. Deep domain generalization via conditional invariant adversarial networks. In *ECCV*, 2018.
- [23] M. Mancini, Z. Akata, E. Ricci, and B. Caputo. Towards recognizing unseen categories in unseen domains. In *ECCV*, 2020.
- [24] U. Maniyar, K. J. Joseph, A. Deshmukh, Ü. Dogan, and V. Balasubramanian. Zero-shot domain generalization. *ArXiv*, 2020.
- [25] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. In *ICLR*, 2013.
- [26] A. Mishra, S. K. Reddy, A. Mittal, and H. A. Murthy. A generative model for zero shot learning using conditional variational autoencoders. *CVPRW*, 2018.
- [27] K. Muandet, D. Balduzzi, and B. Schölkopf. Domain generalization via invariant feature representation. *ArXiv*, abs/1301.2115, 2013.
- [28] S. Narayan, A. Gupta, F.S. Khan, C.G.M. Snoek, and L. Shao. Latent embedding feedback and discriminative features for zero-shot classification. In *ECCV*, 2020.
- [29] T. Nguyen, T. Le, H. Vu, and D. Q. Phung. Dual discriminator generative adversarial nets. In *NeurIPS*, 2017.

- [30] J. Ni, S. Zhang, and H. Xie. Dual adversarial semantics-consistent network for generalized zero-shot learning. *NeurIPS*, 2019.
- [31] X. Peng, Q. Bai, X. Xia, Z. Huang, K. Saenko, and B. Wang. Moment matching for multi-source domain adaptation. (*ICCV*), pages 1406–1415, 2019.
- [32] B. Romera-Paredes and P. H. Torr. An embarrassingly simple approach to zero-shot learning. *ICML*, 2015.
- [33] E. Schonfeld, S. Ebrahimi, S. Sinha, T. Darrell, and Z. Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. *CVPR*, 2019.
- [34] S. Shankar, V. Piratla, S. Chakrabarti, S. Chaudhuri, P. Jyothi, and S. Sarawagi. Generalizing across domains via cross-gradient training. *ArXiv*, 2018.
- [35] R. Socher, M. Ganjoo, C.D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. *NeurIPS*, 2013.
- [36] Y.-H. H. Tsai, L.-K. Huang, and R. Salakhutdinov. Learning robust visual-semantic embeddings. *ICCV*, 2017.
- [37] R. Volpi, H. Namkoong, O. Sener, J. C. Duchi, V. Murino, and S. Savarese. Generalizing to unseen domains via adversarial data augmentation. In *NeurIPS*, 2018.
- [38] M. R. Vyas, H. Venkateswara, and S. Panchanathan. Leveraging seen and unseen semantic relationships for generative zero-shot learning. In *ECCV*, 2020.
- [39] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016.
- [40] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata. Feature generating networks for zero-shot learning. *CVPR*, 2018.
- [41] Y. Xian, S.Choudhury, Y. He, B. Schiele, and Z. Akata. Semantic projection network for zero-and few-label semantic segmentation. In *CVPR*, pages 8256–8265, 2019.
- [42] Y. Xian, S. Sharma, B. Schiele, and Z. Akata. F-vaegan-d2: A feature generating framework for any-shot learning. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10267–10276, 2019.
- [43] Z. Xu, W. Li, L. Niu, and D. Xu. Exploiting low-rank structure from latent domains for domain generalization. In *ECCV*, 2014.
- [44] P. Yang and W. Gao. Multi-view discriminant transfer learning. In *IJCAI*, 2013.