

# Task-Independent Knowledge Makes for Transferable Representations for Generalized Zero-Shot Learning

Chaoqun Wang<sup>1 2</sup>, Xuejin Chen<sup>\*1 2</sup>, Shaobo Min<sup>2</sup>, Xiaoyan Sun<sup>2</sup>, Houqiang Li<sup>1 2</sup>

<sup>1</sup>School of Data Science

<sup>2</sup>The National Engineering Laboratory for Brain-inspired Intelligence Technology and Application  
University of Science and Technology of China, Hefei, Anhui, China

cq14@mail.ustc.edu.cn, xjchen99@ustc.edu.cn, mbobo@mail.ustc.edu.cn, {sunxiaoyan, lihq}@ustc.edu.cn

## Abstract

Generalized Zero-Shot Learning (GZSL) targets recognizing new categories by learning transferable image representations. Existing methods find that, by aligning image representations with corresponding semantic labels, the semantic-aligned representations can be transferred to unseen categories. However, supervised by only seen category labels, the learned semantic knowledge is highly task-specific, which makes image representations biased towards seen categories. In this paper, we propose a novel Dual-Contrastive Embedding Network (DCEN) that simultaneously learns task-specific and task-independent knowledge via semantic alignment and instance discrimination. First, DCEN leverages task labels to cluster representations of the same semantic category by cross-modal contrastive learning and exploring semantic-visual complementarity. Besides task-specific knowledge, DCEN then introduces task-independent knowledge by attracting representations of different views of the same image and repelling representations of different images. Compared to high-level seen category supervision, this instance discrimination supervision encourages DCEN to capture low-level visual knowledge, which is less biased toward seen categories and alleviates the representation bias. Consequently, the task-specific and task-independent knowledge jointly make for transferable representations of DCEN, which obtains averaged 4.1% improvement on four public benchmarks.

## Introduction

Deep learning-based methods are highly successful on various computer vision tasks, such as image classification (He et al. 2016), object detection (Girshick et al. 2014), and semantic segmentation (Badrinarayanan, Kendall, and Cipolla 2017). However, these models require a huge demand for manually labelled training data for numerous classes. To this end, Generalized Zero-Shot Learning (GZSL), which aims to recognize either seen or unseen categories thereby reduces manual annotation labors, recently has attracted great interests.

Due to unavailable unseen category data during training, GZSL puts a high demand on transferability of image

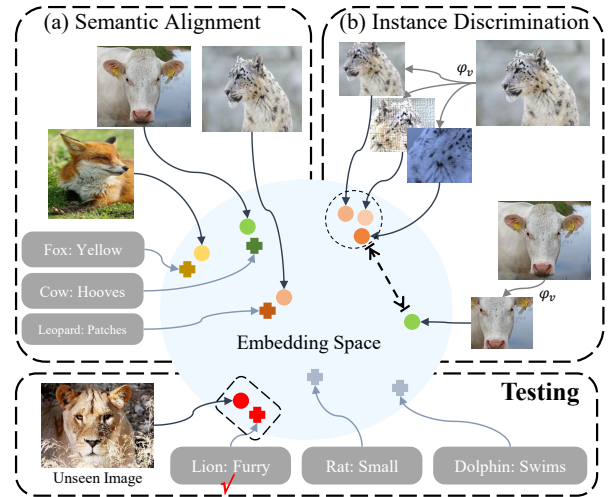


Figure 1: Motivation of this paper. (a) Existing methods focus on using task labels to learn semantic-aligned representations, which can be transferred to unseen categories. (b) Besides, this paper further learns task-independent knowledge via instance discrimination supervision, which significantly improves the representation transferability.

representations (Jiang et al. 2018; Tong et al. 2019). Previous methods tackle this problem by aligning image representations with corresponding category attributes (Xian et al. 2018a), as shown in Fig. 1 (a). Since seen and unseen domain categories share a common attribute space, the semantic-aligned representation can be transferred to unseen categories, and the recognition becomes a nearest neighbour searching problem, *i.e.*, images and category attributes serve as queries and anchors, respectively. Based on this paradigm, most GZSL methods focus on designing elaborate semantic-visual alignment mechanism, such as visual and attribute attentions (Zhu et al. 2019; Xie et al. 2019; Huynh and Elhamifar 2020), semantic auto-encoders (Felix et al. 2018; Min et al. 2019), and domain detectors (Atzmon and Chechik 2019). However, trained with only seen category labels, the learned image representations are highly task-specific and biased towards seen categories, which makes an unseen image tend to be recognized as seen categories.

To alleviate this issue, we introduce task-independent

\*Corresponding Author

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

knowledge that is learned without seen category concepts and shared by both seen and unseen images. We propose a novel Dual-Contrastive Embedding Network (DCEN) that simultaneously learns task-specific and task-independent knowledge to obtain transferable representations in GZSL. Specifically, we design two modules in DCEN, *i.e.*, Semantic Contrastive Module (SCM) and Visual Contrastive Module (VCM), to enhance the semantic-visual alignment knowledge via task labels and introduce annotation-free visual knowledge via instance discrimination, respectively.

First, SCM is developed to leverage semantic labels to learn task-specific knowledge. By constructing cross-modal triplets, SCM constrains the image representation to be not only aligned with the corresponding category but also far away from the most confusing category, which improves the inter-class representation discrimination. Besides, we design masked attribute prediction to complete the missing category attributes from corresponding image representations. This renders the image representation to preserve the semantic-visual complementarity, resulting in better semantic-visual alignment. With cross-modal contrastive learning and masked attribute prediction, SCM can better bridge the semantic-visual gap than previous methods to learn semantic-aligned representation.

Second, VCM is designed to learn task-independent knowledge by distinguishing each image as an individual category, as shown in Fig. 1 (b). This forces the image representation to capture as many detailed visual cues as possible for each individual image, so that all images can be separated apart. Without task labels, the learned low-level visual knowledge is only related to intra-image invariance and inter-image difference, which is irrelevant to seen category concepts. Thus, the learned knowledge is less biased towards the seen categories than the task-specific knowledge from seen category labels, and thereby significantly improves the representation transferability.

Consequently, the task-independent knowledge from instance discrimination and task-specific knowledge from semantic alignment jointly make for transferable representations of our DCEN, which achieves averaged 4.1% improvement on four widely-used GZSL benchmarks. Our overall contribution is threefold:

- We propose a novel Dual-Contrastive Embedding Network (DCEN) that simultaneously learns task-specific and task-independent knowledge to produce transferable representations in GZSL. To the best of our knowledge, this is the first work that introduces task-independent knowledge from instance discrimination to boost representation transferability in GZSL.
- DCEN learns task-independent knowledge by exploring inter-image visual distinction and intra-image transformation invariance. Extensive experiments are conducted to search important low-level invariance concepts for GZSL.
- DCEN learns task-specific knowledge by constructing cross-modal contrastive learning and exploring semantic-visual complementarity to better bridge the semantic-visual gap.

## Related Works

### Generalized Zero-Shot Learning

This paper belongs to Inductive GZSL (Xian et al. 2016; Kodirov, Xiang, and Gong 2017; Min et al. 2020), where the unseen domain data is unavailable during training. A general solution is to learn a joint embedding space, where the image representations and semantic labels, *e.g.*, category attributes (Farhadi et al. 2009; Morgado and Vasconcelos 2017) or text descriptions (Ba et al. 2015), are aligned (Xian et al. 2018a; Zhu, Wang, and Saligrama 2019).

Based on this motivation, most existing methods focus on boosting the semantic-visual alignment by carefully designing the embedding space. For example, previous methods (Shigeto et al. 2015; Zhu, Wang, and Saligrama 2019; Min et al. 2019) constrain the embedding space to be spanned by high-dimensional visual representations, instead of semantic labels, to improve the space discrimination. This can prevent a few seen categories from being the anchors for most input images (Annadani and Biswas 2018), but it also weakens the semantic relationship in the embedding space. To this end, recent methods (Kodirov, Xiang, and Gong 2017; Tong et al. 2019) use auto-encoders to enhance the semantic relationship in the embedding space. For example, DSEN (Min et al. 2019) designs domain-specific cyclic encoders between semantic embeddings and category attributes to explore semantic domain differences, and SP-AEN (Chen et al. 2018) uses cyclic constraints between visual representations and images to capture detailed visual clues. Recently, attention mechanisms are leveraged for images and semantic labels to improve representation discrimination. In VSE (Zhu, Wang, and Saligrama 2019) and SGMA (Zhu et al. 2019), multi-head attentions are designed to infer important local regions to boost semantic-visual alignment. LFGAA (Liu et al. 2019), and DAZLE (Huynh and Elhamifar 2020) use attribute attention to generate clean semantic embeddings. Although being effective, the previous methods focus on better leveraging seen domain labels to improve semantic-visual alignment. Thus, it inevitably makes the image representation highly task-specific and biased towards seen categories, which limits the transferability to unseen categories.

### Self-Supervised Learning

Without human annotations, self-supervised learning methods (Tian, Krishnan, and Isola 2020) have proved that learning task-independent knowledge can produce representations with strong generalization on various downstream tasks, which even surpass fully-supervised methods. Specifically, they usually construct some proxy tasks as supervision to learn image representations. For example, RotNet (Gidaris, Singh, and Komodakis 2018) constrains the model to predict the rotation angles of an input image. CFN (Noroozi and Favaro 2016) regards an image as a jigsaw and predicts the shuffled patch order. Deep clustering methods (Yang et al. 2018; Caron et al. 2018; Yang et al. 2020) separate different classes via clustering. Recently, contrastive learning (Tian, Krishnan, and Isola 2020; He et al. 2020) has attracted increasing attention, which aims to minimize the mutual information between different images and re-

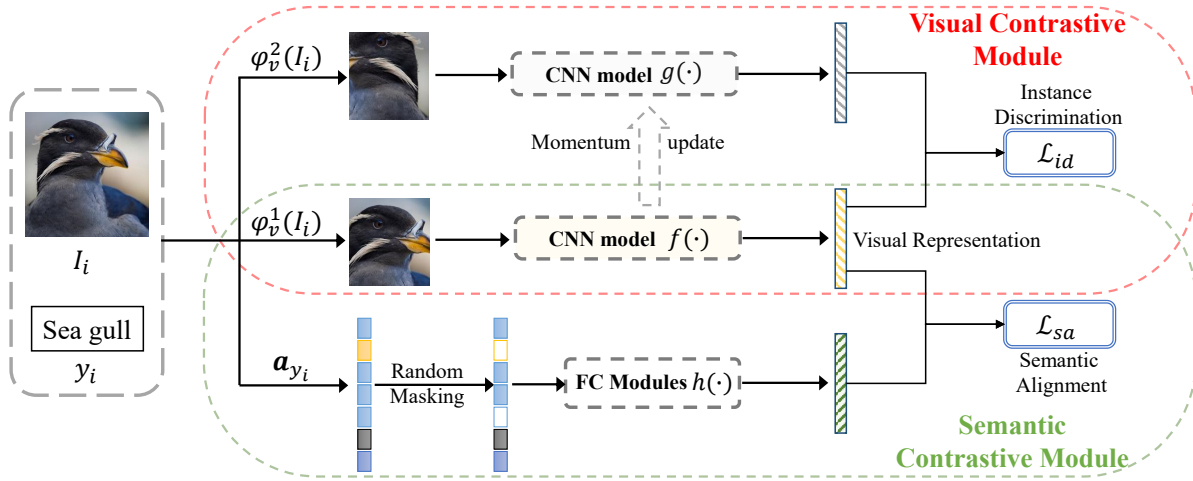


Figure 2: The diagram of Dual-Contrastive Embedding Network. Visual Contrastive Module learns task-independent knowledge via instance discrimination, and Semantic Contrastive Module learns task-specific knowledge via semantic alignment.

gards each image as an individual category. Some works extend contrastive learning to cross-modal tasks. AVID (Morgado, Vasconcelos, and Misra 2020) utilizes the correspondence between audio and visual representations. CMC (Tian, Krishnan, and Isola 2020) contrasts across multiple views, such as luminance, chrominance RGB, and depth. Inspired by the generalized self-supervised representation, task-independent knowledge has been introduced in many other fields, such as few-shot learning (Gidaris et al. 2019) and landmark detection (Cheng, Su, and Maji 2020).

In this paper, we explore the task-independent knowledge from instance discrimination to boost representation transferability in GZSL. Notably, SDGN (Wu et al. 2020) also uses self-supervised learning in GZSL. However, it is much different from our DCEN, because: (a) SDGN follows Transductive ZSL setting that uses unseen domain images for training, while our DCEN only uses data from seen domain during the training stage; (b) SDGN applies contrastive loss to triplet data of two domains to alleviate domain confusion, while we utilize instance discrimination knowledge of the seen domain; (c) SDGN is a generative method that utilizes powerful GANs for unseen domain feature generation.

## Dual-Contrastive Embedding Network

### Problem Formulation

The target of GZSL is to recognize images from either seen or unseen categories using the model trained only with seen domain data. Here, we define seen domain data as  $\mathcal{S} = \{I_i, y_i, \mathbf{a}_{y_i} | I_i \in \mathcal{I}_s, y_i \in \mathcal{Y}_s, \mathbf{a}_{y_i} \in \mathcal{A}_s\}$ , where  $I_i$  is an image of a seen category,  $y_i$  is the corresponding category, and  $\mathbf{a}_{y_i}$  is the semantic description of  $y_i$ , such as category attributes. Unseen domain data are similarly defined as  $\mathcal{U} = \{I_i, y_i, \mathbf{a}_{y_i} | I_i \in \mathcal{I}_u, y_i \in \mathcal{Y}_u, \mathbf{a}_{y_i} \in \mathcal{A}_u\}$ , where  $\mathcal{Y}_s \cap \mathcal{Y}_u = \emptyset$ .

One main stream of GZSL methods is to learn a joint embedding space, where the image representations and cate-

gory descriptions can be aligned by minimizing:

$$\mathcal{L}_{zsl} = \sum_{I_i \in \mathcal{I}_s} d(f(I_i), \mathbf{a}_{y_i}), \quad (1)$$

where  $f(\cdot)$  is the feature extractor to produce image representations.  $d(\cdot, \cdot)$  is a distance function, such as negative cosine similarity. Since  $\mathcal{A}_s$  and  $\mathcal{A}_u$  usually share a common semantic space, the semantic-visual relationship captured by  $f(\cdot)$  can be transferred to the unseen domain by:

$$\hat{y} = \arg \min_{y \in \mathcal{Y}_s \cup \mathcal{Y}_u} d(f(I), \mathbf{a}_y), \quad (2)$$

where  $I \in \mathcal{I}_s \cup \mathcal{I}_u$ .

However, due to unavailable  $\mathcal{U}$  during training, it puts a high demand on transferability of learned  $f(\cdot)$  in Eq. (1).

### Semantic Contrastive Module

In order to better align semantic and visual embeddings, we introduce a Semantic Contrastive Module (SCM) to better bridge the semantic-visual gap via cross-modal contrastive learning and exploring semantic-visual complementarity. Specifically, in the SCM, we construct semantic-visual triplets to prevent image representations from being biased towards confusing category and design masked attribute prediction to complete the missing attributes from image representations.

As shown in Fig. 2, given an image  $I_i$  and its category attribute  $\mathbf{a}_{y_i}$ , SCM learns the semantic-aligned representation  $f(I_i)$  by minimizing:

$$\mathcal{L}_{sa} = \sum_{I_i \in \mathcal{I}_s} d(f(I_i), h(\hat{\mathbf{a}}_{y_i})) - \min_{y_k \neq y_i} \{d(f(I_i), h(\hat{\mathbf{a}}_{y_k}))\}. \quad (3)$$

Different from the general objective function Eq. (1), SCM first randomly masks some elements of  $\mathbf{a}_{y_i}$ , which generates  $\hat{\mathbf{a}}_{y_i}$  to introduce some semantic variance and improve model robustness to input noises. As attributes contain abstract semantic information, such as wing colors, it is hard for visual feature extractor  $f(\cdot)$  to extract such highly abstract




Figure 3: Architecture for masked attribute prediction.

concepts from the images. Thus, instead of directly aligning  $f(I_i)$  with  $\hat{a}_{y_i}$ , SCM uses  $K$  fully-connected (FC) modules, consisting of FC, Batch normalization (BN), and ReLU, to produce high-level semantic embeddings  $h(\hat{a}_{y_i})$ , which can better bridge the semantic-visual gap. Besides, SCM constructs cross-modal triplet  $\{f(I_i), h(\hat{a}_{y_i}), h(\hat{a}_{y_k})\}$ , where  $f(I_i)$  is the anchor image representation,  $h(\hat{a}_{y_i})$  is the positive category semantic embedding, and  $h(\hat{a}_{y_k})$  is the most confusing negative category semantic embedding. Triplet loss has been demonstrated useful in previous works of face recognition (Schroff, Kalenichenko, and Philbin 2015). Thus,  $\mathcal{L}_{sa}$  can not only align  $f(I_i)$  with the corresponding  $h(\hat{a}_{y_i})$ , but also punish the most confusing category. Compared to  $\mathcal{L}_{zsl}$ ,  $\mathcal{L}_{sa}$  can produce more discriminative semantic-aligned  $f(I_i)$  via cross-modal contrastive learning.

To better preserve semantic relationships among categories and encode the semantic knowledge into  $f(I_i)$ , masked attribute prediction  $\hat{h}(\cdot)$  is designed by predicting the missed elements in  $\hat{a}_{y_i}$  from the image representation  $f(I_i)$ . The architecture is shown in Fig. 3. we first fuse the image representation  $f(I_i)$  and semantic embedding  $h(\hat{a}_{y_i})$  via FC blocks and concatenation. Then, a decoder module, consisting of  $K$  FC blocks, is used to predict the intact attribute  $\mathbf{a}_{y_i}^p$ . The depths of  $h(\cdot)$  and  $\hat{h}(\cdot)$  are kept consistent for better attribute reconstruction. Notably, since the attribute masking is random and  $h(\hat{a}_{y_i})$  has no masking knowledge,  $f(I_i)$  should capture accurate and entire semantic information from  $I_i$  to complete  $\hat{a}_{y_i}$ , thereby enhancing the semantic-visual alignment. The loss function for masked attribute prediction is:

$$\mathcal{L}_{sp} = \sum_{I_i \in \mathcal{I}_s} \left\| \hat{h}(f(I_i), h(\hat{a}_{y_i})) - \mathbf{a}_{y_i} \right\|_2. \quad (4)$$

With cross-modal contrastive learning and masked attribute prediction, SCM learns better semantic alignment knowledge than naive  $\mathcal{L}_{zsl}$ , which makes image representation  $f(\cdot)$  better transferable to unseen categories.

### Visual Contrastive Module

Since only  $\mathcal{S}$  is available for training, the learned semantic knowledge in  $f(\cdot)$  from SCM is inevitably task-specific to the seen categories. Thus, a visual contrastive module (VCM) is proposed to learn task-independent knowledge without human annotations, which further improves the transferability of  $f(\cdot)$ . To this end, VCM regards each im-

age  $I$  as an individual category, to push different image representations far away from each other. This enforces  $f(\cdot)$  to capture more detailed visual clues to represent each individual image so that they can be separated apart well. Besides, VCM learns low-level visual invariance concepts, such as rotation and cropping invariance, that lead to strong generalization.

Given the  $i$ -th image  $I_i$  in  $\mathcal{I}_s$ , VCM first transforms  $I_i$  into two different counterparts  $I_i^{v1} = \varphi_v^1(I_i)$  and  $I_i^{v2} = \varphi_v^2(I_i)$ .  $\varphi_v(\cdot)$  is a composition of random low-level transformations, e.g., cropping, rotation, blurring, or color jittering. Then,  $I_i^{v1}$  and  $I_i^{v2}$  are fed into two CNN models  $f(\cdot)$  and  $g(\cdot)$  to extract visual representations  $f(I_i^{v1})$  and  $g(I_i^{v2})$ , respectively. Finally, a visual contrastive loss is used to attract representations of two counterparts generated from the same image and repel representations from different images by:

$$\mathcal{L}_{id} = - \sum_{I_i \in \mathcal{I}_s} \log \frac{\exp\{d(f(I_i^{v1}), g(I_i^{v2}))/\tau\}}{\sum_j \exp\{d(f(I_i^{v1}), g(I_j^{v2}))/\tau\}}, \quad (5)$$

where  $d(\cdot, \cdot)$  is the cosine similarity between two representations, and  $\tau$  is a temperature parameter. The positive data pair is defined as  $\{I_i^{v1}, I_i^{v2}\}$  that are two views from the same image, while the negative data pair is defined as  $\{I_i^{v1}, I_j^{v2}\}$  for different images when  $i \neq j$ . Data from previous batches serve as negative data. The numerator term of  $\mathcal{L}_{id}$  constrains  $f(\cdot)$  to be invariant to pre-defined image transformations  $\varphi_v(\cdot)$ . By designing different image transformations  $\varphi_v(\cdot)$ , VCM can induce different low-level invariance concepts that are useful in GZSL tasks. The denominator term of  $\mathcal{L}_{id}$  constrains  $f(\cdot)$  to be discriminative among different images.

In  $\mathcal{L}_{id}$ , the weights of  $g(\cdot)$  are momentum-updated (Laine and Aila 2017) from  $f(\cdot)$  by:

$$W_g = m * W_g + (1 - m) * W_f, \quad (6)$$

where  $W_g$  and  $W_f$  are weights of  $g(\cdot)$  and  $f(\cdot)$ .  $m$  is a momentum parameter. The momentum-updated  $g(\cdot)$  has two main advantages: a) compared to a new trainable CNN,  $g(\cdot)$  has less memory cost, and b) compared to sharing weights of  $f(\cdot)$ ,  $g(\cdot)$  is the temporal ensemble of different checkpoints from  $f(\cdot)$  and contains temporal momentum information.

Consequently, through instance discrimination  $\mathcal{L}_{id}$ ,  $f(\cdot)$  can gather augmented views of the same image and push different images apart. Thus,  $f(\cdot)$  can not only capture the unique information contained in each individual image but also learn useful invariance concepts. More importantly, this low-level visual knowledge is irrelevant to any category concept of  $\mathcal{Y}_s$ , thereby being well-transferable to the unseen images.

### Overall Objective

Finally, the overall objective function of DCEN is :

$$\mathcal{L}_{all} = \lambda_1 \mathcal{L}_{id} + \mathcal{L}_{sa} + \lambda_2 \mathcal{L}_{sp}, \quad (7)$$

where  $\lambda_1$  and  $\lambda_2$  are hyper-parameters to balance  $\mathcal{L}_{id}$  and  $\mathcal{L}_{sp}$ . The task-specific knowledge of  $\mathcal{L}_{sa} + \lambda_2 \mathcal{L}_{sp}$  and task-independent knowledge of  $\lambda_1 \mathcal{L}_{id}$  jointly make for transferable visual representation  $f(I)$ .

Dataset	Seen/Unseen	Attributes	Train	Val	Test
CUB	150/50	312	7,057	1,764	2,967
AWA2	40/10	85	23,527	5,882	7,913
aPY	20/12	64	5,932	1,483	7,924
SUN	645/72	102	10,320	2,580	1,440

Table 1: Detailed statistics of datasets.




Figure 4: Effects of different  $\lambda_1$  and  $\lambda_2$  for  $\mathcal{L}_{id}$  and  $\mathcal{L}_{sp}$ , respectively.

The testing phase is viewed as a nearest neighbour searching process by:

$$\hat{y} = \arg \min_{y \in \mathcal{Y}_s \cup \mathcal{Y}_u} d(f(I), h(\mathbf{a}_y)), \quad (8)$$

where  $I \in \mathcal{I}_s \cup \mathcal{I}_u$  is the query image, and  $\mathbf{a}_y$  is the anchor category attribute from both the seen and unseen domains.

## Experiments

In this section, we first evaluate each component of DCEN and then compare DCEN with state-of-the-art GZSL methods. Please refer to the supplementary for implementation details of DCEN.

### Experimental Settings

**Datasets.** We adopt four widely-used GZSL benchmarks, which are Caltech-USCD Birds-200-2011 (CUB) (Wah et al. 2011), SUN (Patterson and Hays 2012), Animals with Attributes2 (AWA2) (Xian et al. 2018a), and Attribute Pascal and Yahoo (aPY) (Farhadi et al. 2009) for the following experiments. Detailed statistics of datasets are listed in Table 1. Category attributes provided by datasets are used as semantic labels.

**Metrics.** For GZSL, the harmonic mean  $H = (2MCA_u \times MCA_s) / (MCA_u + MCA_s)$  is widely used to evaluate the model performance, where  $MCA_s$  and  $MCA_u$  are the Mean Class Top-1 Accuracy for seen and unseen domains, respectively.

### Ablation Study

In this part, we analyze each component of DCEN on CUB and aPY datasets for fast validation. The baseline model is denoted as Basic-ZSL, which only learns  $f(\cdot)$  via  $\mathcal{L}_{ZSL}$  in Eq. (1) and is test via Eq. (2).

**Effects of instance discrimination.** The task-independent knowledge from instance discrimination is a major contribution of this paper, which is learned by Visual Contrastive Module (VCM). Thus, we first analyze how the task-independent knowledge boosts the transferability of

Aug. prob.	Crop 1.0	Flip 0.5	Gray 0.2	CJ 0.8	Blur 0.5	Rotation 0.5	Swap 0.2	H
								65.3
	✓							65.8↑
	✓	✓						66.0↑
	✓	✓	✓					65.5↓
	✓	✓		✓(v1)				63.8↓
	✓	✓		✓(v2)				65.3↓
	✓	✓		✓(v3)				65.2↓
	✓	✓			✓			67.1↑
$\mathcal{L}_{id}$	✓	✓				✓(90)		64.4↓
	✓	✓				✓(60)		64.7↓
	✓	✓				✓(30)		66.5↑
	✓	✓					✓(7)	64.2↓
	✓	✓					✓(5)	65.8↓
	✓	✓					✓(3)	66.7↑
	✓	✓			✓	✓(30)		67.9↑
	✓	✓			✓	✓(30)	✓(3)	<b>68.5↑</b>

Table 2: Evaluating different visual augmentations on CUB by successively adding operations. When a certain operation brings positive effects, it is retained, otherwise, it is removed. The operation probability is experimentally determined. For v1, v2, and v3 color jittering, the parameters are respectively (0.4, 0.4, 0.4, 0.4), (0.4, 0.4, 0.4, 0.1), and (0.8, 0.8, 0.8, 0.2) for brightness, contrast, saturation, and hue. For rotation and swap, the parameters in parentheses indicate rotation angle and jigsaw number, respectively.

semantic-aligned representation in GZSL by varying  $\lambda_1$  to influence the visual contrastive loss  $\mathcal{L}_{id}$ . The results are given in Fig. 4 (a). In this experiment, we add VCM to Basic-ZSL with varying  $\lambda_1$ , thus  $\lambda_1 = 0$  denotes no task-independent knowledge introduced. We can see that when increasing  $\lambda_1$  from 0 to 0.1 for CUB and aPY,  $\mathcal{L}_{id}$  begins introducing instance discriminative knowledge, and  $H$  gradually increases. The best recognition performance is obtained at around  $\lambda_1 = 0.1$  for both CUB and aPY. This demonstrates that the task-independent knowledge brought by instance discrimination  $\mathcal{L}_{id}$  enables DCEN to produce more transferable representations, thereby recognizing unseen domain images more accurately. When  $\lambda_1$  increases from 0.1 to 1.0, the performance  $H$  drops. Specifically, for CUB, the result with  $\lambda_1 = 1$  is even worse than that with  $\lambda_1 = 0$ . This reveals that, when  $\lambda_1$  is large, the instance discrimination knowledge  $\mathcal{L}_{id}$ , which targets pushing all representations far away, becomes too sensitive to intra-class variance of different images and hard to preserve the semantic relationship, thereby being unable to cluster images correctly. Thus, a trade-off between task-independent knowledge and task-specific knowledge is critical to make for strong transferable representations in GZSL. We find  $\lambda_1 = 0.1$  is suitable for most cases, which is used for the following experiments.

**Effects of different invariance concepts.** In the task-independent knowledge brought by instance discrimination, the low-level invariance concepts, induced by visual augmentations, are important. In this part, we conduct extensive experiments to explore useful invariance concepts in GZSL by combining different augmentations as  $\varphi_v(\cdot)$  in VCM on CUB. The experiments are carried out based on Basic-ZSL with  $\lambda_1 = 0.1$ . The results are listed in Table 2, where we explore seven widely-used low-level augmenta-

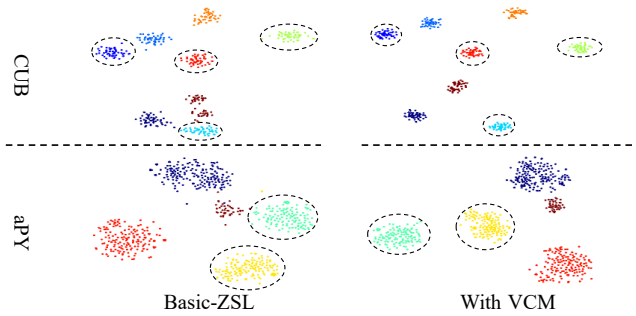


Figure 5: Visualization of representation on CUB and aPY. Dotted circles denote unseen domain representations.

tions, which contain random cropping, horizontal flipping, grayscale, color jittering, Gaussian blurring, rotation, and swapping. Specifically, the swapping augmentation (Chen et al. 2019) regards an image as a jigsaw and randomly shuffles the patches, which requires the representation to be invariant to patch-swapped images. From the results, adding cropping, horizontal flipping, Gaussian blurring, rotation, and swap augmentations into  $\varphi_v(\cdot)$  can benefit the performance, while adding grayscale and color jittering harms the results. This is because that the grayscale and color augmentations will change the color space severely, which may prevent the model from inferring correct color attributes, e.g., wing color, to recognize unseen categories. Among the useful augmentations, Gaussian blurring brings the largest improvement because it helps the model neglect high-frequency noises. For random rotation and swapping, too large angles and jigsaw numbers bring negative effects, which may destroy the image contents severely. Finally, we combine all positive low-level augmentations in Table 2 as  $\varphi_v(\cdot)$ , which obtains an impressive improvement of 3.2% over the baseline. The useful low-level invariance concepts explored in this experiment may provide an insight for the following GZSL studies.

**Visualization of representation.** To intuitively analyze the effect of task-independent knowledge from instance discrimination, we visualize the image representations before and after adding VCM using t-SNE (Van der Maaten and Hinton 2008) on randomly selected categories. The results are given in Fig. 5. After introducing task-independent knowledge into image representations, the category margin is obviously enlarged between seen and unseen categories, and the distribution of each individual cluster becomes more compact. This means that image representations are more distinguishable between categories of the two domains, thus the unseen image representation is less biased towards the seen categories. In other words, the visual representation can be better transferred to represent unseen image characteristics. Besides, Table 3 gives quantitative results after adding VCM. It shows that VCM can bring an impressive gain of 5% on  $H$  and 6.1% on  $MCA_u$ , which demonstrates that VCM helps the model better transfer to the unseen domain. In summary, both quantitative and qualitative results prove that the task-independent knowledge from instance discrimination can improve the representation transferabil-

Methods	$MCA_u$	$MCA_s$	$H$
Basic-ZSL	56.3	72.8	63.5
Basic-ZSL with VCM	62.4	75.9	68.5
+attribute masking	62.5	78.3	69.5
+cross-modal triplet	63.5	77.7	69.9
+ $\mathcal{L}_{sp}$	63.8	78.4	<b>70.4</b>

Table 3: Effects of each component of DCEN on CUB.




Figure 6: Effects of different  $\sigma$  and  $K$ .

ity in GZSL.

**Analysis of semantic contrastive module.** In this part, we analyze each component of SCM, and the results are given in Table 3. Based on Basic-ZSL with VCM, we add random attribute masking, cross-modal triplet loss, and masked attribute prediction sequentially. First, we evaluate the effect of masking  $\sigma\%$  attributes in SCM, which can improve model robustness to semantic noises. In training, attributes are randomly chosen with  $p = 0.5$ , and  $\sigma\%$  elements of these chosen attributes are masked. Random masking is not used in testing. Fig. 6 (a) shows the effect of varying  $\sigma$  on CUB and aPY. When  $\sigma$  is too large which means many elements of attributes are missed, the semantic alignment becomes weak, thereby the performance drops. Notably, the reason for small optimal  $\sigma$  on aPY is that there are much fewer attributes in aPY than in CUB. Thus, suitable semantic masking is a simple way to improve semantic-visual robustness. Then, adding cross-modal triplet loss  $\mathcal{L}_{sa}$  brings about 0.4% improvement on CUB in Table 3. Specifically,  $\mathcal{L}_{sa}$  contrasts cross-modal triplets to make visual representation not only close to corresponding semantic anchors but also far away from confusing negative categories, which improves representation discrimination. Finally, we evaluate the effect of masked attribute prediction by varying  $\lambda_2$  of  $\mathcal{L}_{sp}$ , and the results are given in Fig. 4 (b). We can see that setting appropriate weights of  $\lambda_2$  for  $\mathcal{L}_{sp}$  can boost the final results. It proves that, by constraining the visual representation to predict the missed attributes, more semantic information from category attributes is explicitly encoded and preserved in visual representation. In summary, our SCM can generate better semantic-aligned representation via cross-modal triplet loss and masked attribute prediction.

**Analysis for other hyper-parameters.** In DCEN,  $K$ ,  $\tau$ , and  $m$  are three minor hyper-parameters.  $K$  determines the architecture  $h(\cdot)$  in Eq. (3). Here, we evaluate the effects of different  $K$  in Fig. 6 (b). It can be seen that  $K = 2$  is suitable for most cases.  $\tau$  and  $m$  control the cosine similarity scaling in Eq. (5) and momentum updating of  $g(\cdot)$  in Eq. (6). In this paper, we set  $\tau = 0.07$  and  $m = 0.999$ , which are

Methods	CUB			AWA2			aPY			SUN			
	$MCA_u$	$MCA_s$	$H$	$MCA_u$	$MCA_s$	$H$	$MCA_u$	$MCA_s$	$H$	$MCA_u$	$MCA_s$	$H$	
GEN	FGN(Xian et al. 2018b)	43.7	57.7	49.7	-	-	-	-	-	42.6	36.6	39.4	
	SABR-I(Paul, Krishnan, and Munjal 2019)	55.0	58.7	56.8	30.3	93.9	46.9	-	-	50.7	35.1	41.5	
	f-VAEGAN-D2(Xian et al. 2019)	63.2	75.6	68.9	-	-	-	-	-	50.1	37.8	43.1	
	RFF-GZSL(Han, Fu, and Yang 2020)	59.8	79.9	68.4	-	-	-	-	-	58.8	45.3	51.2	
	OCD-CVAE(Keshari, Singh, and Vatsa 2020)	44.8	59.9	51.3	59.5	73.4	65.7	-	-	44.8	42.9	43.8	
	IZF-Softmax(Shen, Qin, and Huang 2020)	52.7	68.0	59.4	60.6	77.5	68.0	42.3	60.5	49.8	52.7	57.0	54.8
	TF-VAEGAN(Narayan et al. 2020)	63.8	79.3	70.7	-	-	-	-	-	62.4	47.1	53.7	
NON-GEN	SYNC(Changpinyo et al. 2016)	11.5	70.9	19.8	10.0	90.5	18.0	7.4	66.3	13.3	7.9	43.3	13.4
	CDL(Jiang et al. 2018)	23.5	55.2	32.9	-	-	-	19.8	48.6	28.1	21.5	34.7	26.5
	PSR-ZSL(Annadani and Biswas 2018)	24.6	54.3	33.9	20.7	73.8	32.2	13.5	51.4	21.4	20.8	37.2	26.7
	SP-AEN(Chen et al. 2018)	34.7	70.6	46.6	23.3	90.9	37.1	13.7	63.4	22.6	24.9	38.6	30.3
	DLFZRL(Tong et al. 2019)	-	-	37.1	-	-	45.1	-	-	31.0	-	-	24.6
	MLSE(Ding and Liu 2019)	22.3	71.6	34.0	23.8	83.2	37.0	12.7	74.3	21.7	20.7	36.4	26.4
	TripletLoss(Cacheux, Borgne, and Crucianu 2019)	55.8	52.3	53.0	48.5	83.2	61.3	-	-	-	47.9	30.4	36.8
	COSMO(Atzmon and Chechik 2019)	44.4	57.8	50.2	-	-	-	-	-	-	44.9	37.7	41.0
	PREN(Ye and Guo 2019)	32.5	55.8	43.1	32.4	88.6	47.4	-	-	-	35.4	27.2	30.8
	VSE-S(Zhu, Wang, and Saligrama 2019)	33.4	87.5	48.4	41.6	91.3	57.2	24.5	72.0	36.6	-	-	-
	AREN(Xie et al. 2019)	63.2	69.0	66.0	54.7	79.1	64.7	30.0	47.9	36.9	40.3	32.3	35.9
	CosineSoftmax(Kampffmeyer et al. 2019)	47.4	47.6	47.5	56.4	81.4	66.7	26.5	74.0	<u>39.0</u>	36.3	42.8	39.3
	DAZLE(Huynh and Elhamifar 2020)	56.7	59.6	58.1	60.3	75.7	67.1	-	-	-	52.3	24.3	33.2
	<b>DCEN</b>	63.8	78.4	<b>70.4</b>	62.4	81.7	<b>70.8</b>	37.5	61.6	<b>46.7</b>	43.7	39.8	<b>41.7</b>

Table 4: Results of GZSL on four classification benchmarks. Generative methods (GEN) utilize extra synthetic unseen domain data for training. The best result is bolded, and the second best is underlined.

commonly used in contrastive learning (He et al. 2020).

### Comparison with State-of-the-Art Methods

Finally, we compare our DCEN with the state-of-the-art methods on four datasets, *i.e.*, CUB, AWA2, aPY, and SUN. Results are given in Table 4, which shows that DCEN surpasses existing methods on four datasets by a large margin.

First, we compare DCEN with related embedding-based methods, which use only seen domain data for training and no generative models. DCEN surpasses the best related method by respectively 4.4%, 3.7%, 7.7%, 0.7% on CUB, AWA2, aPY, and SUN datasets in the term of  $H$ , demonstrating its good generalizability to different image domains. Notably, among non-generative methods, AREN (Xie et al. 2019) obtains similar performance to DCEN, but it utilizes a model ensemble of two separate recognition branches. Our improvements are impressive, because DCEN introduces extra task-independent knowledge into image representations via instance discrimination learning, and it is proved that appropriate task-independent knowledge can significantly improve representation transferability in GZSL.

Then, we compare DCEN with recent generative methods. Notably, the generative methods utilize prior unseen domain semantics to synthesize extra unseen visual data for training, *e.g.*, powerful GANs, while DCEN only uses the seen domain data. From Table 4, DCEN outperforms most generative methods by a large margin, which is encouraging for embedding-based methods because no synthesized unseen domain data is used for training. Compared to recent TF-VAEGAN (Narayan et al. 2020), DCEN obtains comparable results on CUB using only seen domain data. This is due to introducing task-independent knowledge from instance discrimination, and it proves that the embedding-based methods have much potential.

Finally, we can conclude that: a) task-independent knowledge and task-specific semantic knowledge jointly make for strong transferable representations in GZSL, which enables

DCEN to surpass all related works; and b) task-independent knowledge from instance discrimination is useful, and other task-independent knowledge remains to be explored.

### Discussion

Compared to previous embedding-based methods, the improvement of DCEN on SUN is much less than the other three benchmarks. The reason is that the number of images in each category of SUN is small, *e.g.*, averaged 16 images per category. Thus, learning instance discrimination has a similar effect on image representation, compared to category discrimination. This also reveals a limitation of DCEN, *i.e.*, the task-independent knowledge from instance discrimination may be not suitable for tail categories with inadequate samples.

### Conclusion

In this paper, we propose a novel Dual-Contrastive Embedding Network (DCEN) that utilizes task-independent and task-specific knowledge to jointly make for transferable representation in GZSL. Specifically, a semantic contrastive module is developed to learn task-specific knowledge by performing cross-modal contrastive learning and exploring semantic-visual complementarity with category labels. Besides, a visual contrastive module is designed to learn annotation-free task-independent knowledge via instance discrimination supervision, which gathers representations of the same image and pushes different image representations apart. Compared to seen category knowledge, the task-independent knowledge from instance discrimination is less biased, which can improve the representation transferability to unseen categories. Extensive experiments show that our DCEN achieves superior performance on four GZSL benchmarks.

In the future, the effects of more task-independent knowledge, such as rotation angle prediction and jigsaw order prediction, will be explored for GZSL.

## Acknowledgements

This work was supported by the National Key R&D Program of China with grant No. 2020AAA0108602, National Natural Science Foundation of China (NSFC) under Grants 61632006 and 62076230, and Fundamental Research Funds for the Central Universities under Grants WK3490000003.

## Ethics Statement

All datasets used in this paper are publicly available and have no relations with personal privacy information.

## References

- Annadani, Y.; and Biswas, S. 2018. Preserving Semantic Relations for Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7603–7612.
- Atzmon, Y.; and Chechik, G. 2019. Adaptive Confidence Smoothing for Generalized Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11671–11680.
- Ba, J. L.; Swersky, K.; Fidler, S.; and Salakhutdinov, R. 2015. Predicting Deep Zero-Shot Convolutional Neural Networks Using Textual Descriptions. In *Proceedings of the IEEE International Conference on Computer Vision*, 4247–4255.
- Badrinarayanan, V.; Kendall, A.; and Cipolla, R. 2017. Segnet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39(12): 2481–2495.
- Cacheux, Y. L.; Borgne, H. L.; and Crucianu, M. 2019. Modeling Inter and Intra-Class Relations in the Triplet Loss for Zero-Shot Learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 10333–10342.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep Clustering for Unsupervised Learning of Visual Features. In *Proceedings of the European Conference on Computer Vision*, 132–149.
- Changpinyo, S.; Chao, W.-L.; Gong, B.; and Sha, F. 2016. Synthesized Classifiers for Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5327–5336.
- Chen, L.; Zhang, H.; Xiao, J.; Liu, W.; and Chang, S.-F. 2018. Zero-Shot Visual Recognition using Semantics-Preserving Adversarial Embedding Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1043–1052.
- Chen, Y.; Bai, Y.; Zhang, W.; and Mei, T. 2019. Destruction and Construction Learning for Fine-Grained Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5157–5166.
- Cheng, Z.; Su, J.-C.; and Maji, S. 2020. Unsupervised Discovery of Object Landmarks via Contrastive Learning. *arXiv preprint arXiv:2006.14787*.
- Ding, Z.; and Liu, H. 2019. Marginalized Latent Semantic Encoder for Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6191–6199.
- Farhadi, A.; Endres, I.; Hoiem, D.; and Forsyth, D. 2009. Describing Objects by Their Attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1778–1785.
- Felix, R.; Kumar, V. B.; Reid, I.; and Carneiro, G. 2018. Multi-Modal Cycle-Consistent Generalized Zero-Shot Learning. In *Proceedings of the European Conference on Computer Vision*, 21–37.
- Gidaris, S.; Bursuc, A.; Komodakis, N.; Pérez, P.; and Cord, M. 2019. Boosting Few-Shot Visual Learning with Self-Supervision. In *Proceedings of the IEEE International Conference on Computer Vision*, 8059–8068.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised Representation Learning by Predicting Image Rotations. In *International Conference on Learning Representations*, 1–16.
- Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 580–587.
- Han, Z.; Fu, Z.; and Yang, J. 2020. Learning the Redundancy-Free Features for Generalized Zero-Shot Object Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12865–12874.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Huynh, D.; and Elhamifar, E. 2020. Fine-Grained Generalized Zero-Shot Learning via Dense Attribute-Based Attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4483–4493.
- Jiang, H.; Wang, R.; Shan, S.; and Chen, X. 2018. Learning Class Prototypes via Structure Alignment for Zero-Shot Recognition. In *Proceedings of the European Conference on Computer Vision*, 118–134.
- Kampffmeyer, M.; Chen, Y.; Liang, X.; Wang, H.; Zhang, Y.; and Xing, E. P. 2019. Rethinking Knowledge Graph Propagation for Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11487–11496.
- Keshari, R.; Singh, R.; and Vatsa, M. 2020. Generalized Zero-Shot Learning via Over-Complete Distribution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 13300–13308.
- Kodirov, E.; Xiang, T.; and Gong, S. 2017. Semantic Autoencoder for Zero-Shot Learning. In *Proceedings of the*



- IEEE Conference on Computer Vision and Pattern Recognition*, 3174–3183.
- Laine, S.; and Aila, T. 2017. Temporal Ensembling for Semi-supervised Learning. In *International Conference on Learning Representations*, 1–13.
- Liu, Y.; Guo, J.; Cai, D.; and He, X. 2019. Attribute Attention for Semantic Disambiguation in Zero-Shot Learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 6698–6707.
- Min, S.; Yao, H.; Xie, H.; Zha, Z.-J.; and Zhang, Y. 2019. Domain-Specific Embedding Network for Zero-Shot Recognition. In *Proceedings of ACM International Conference on Multimedia*, 2070–2078.
- Min, S.; Yao, H.; Xie, H.; Zha, Z.-J.; and Zhang, Y. 2020. Domain-Oriented Semantic Embedding for Zero-Shot Learning. *IEEE Transactions on Multimedia*.
- Morgado, P.; and Vasconcelos, N. 2017. Semantically Consistent Regularization for Zero-Shot Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6060–6069.
- Morgado, P.; Vasconcelos, N.; and Misra, I. 2020. Audio-Visual Instance Discrimination with Cross-Modal Agreement. *arXiv preprint arXiv:2004.12943*.
- Narayan, S.; Gupta, A.; Khan, F. S.; Snoek, C. G.; and Shao, L. 2020. Latent Embedding Feedback and Discriminative Features for Zero-Shot Classification. In *Proceedings of the European Conference on Computer Vision*, 1–23.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In *Proceedings of the European Conference on Computer Vision*, 69–84.
- Patterson, G.; and Hays, J. 2012. Sun Attribute Database: Discovering, Annotating, and Recognizing Scene Attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2751–2758.
- Paul, A.; Krishnan, N. C.; and Munjal, P. 2019. Semantically Aligned Bias Reducing Zero Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7056–7065.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. FaceNet: A Unified Embedding for Face Recognition and Clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.
- Shen, Y.; Qin, J.; and Huang, L. 2020. Invertible Zero-Shot Recognition Flows. In *Proceedings of the European Conference on Computer Vision*, 614–631.
- Shigeto, Y.; Suzuki, I.; Hara, K.; Shimbo, M.; and Matsumoto, Y. 2015. Ridge Regression, Hubness, and Zero-Shot Learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 135–151.
- Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive Multiview Coding. In *Proceedings of the European Conference on Computer Vision*, 776–794.
- Tong, B.; Wang, C.; Klinkigt, M.; Kobayashi, Y.; and Nonaka, Y. 2019. Hierarchical Disentanglement of Discriminative Latent Features for Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11467–11476.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9(11): 2579–2605.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The Caltech-UCSD Birds-200-2011 Dataset. In *California Institute of Technology*, 1–8.
- Wu, J.; Zhang, T.; Zha, Z.-J.; Luo, J.; Zhang, Y.; and Wu, F. 2020. Self-Supervised Domain-Aware Generative Network for Generalized Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 12767–12776.
- Xian, Y.; Akata, Z.; Sharma, G.; Nguyen, Q.; Hein, M.; and Schiele, B. 2016. Latent Embeddings for Zero-Shot Classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 69–77.
- Xian, Y.; Lampert, C. H.; Schiele, B.; and Akata, Z. 2018a. Zero-shot learning—A Comprehensive Evaluation of the Good, the Bad and the Ugly. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41(9): 2251–2265.
- Xian, Y.; Lorenz, T.; Schiele, B.; and Akata, Z. 2018b. Feature Generating Networks for Zero-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5542–5551.
- Xian, Y.; Sharma, S.; Schiele, B.; and Akata, Z. 2019. F-VAEGAN-D2: A Feature Generating Framework for Any-Shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 10275–10284.
- Xie, G.-S.; Liu, L.; Jin, X.; Zhu, F.; Zhang, Z.; Qin, J.; Yao, Y.; and Shao, L. 2019. Attentive Region Embedding Network for Zero-shot Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9384–9393.
- Yang, X.; Deng, C.; Liu, X.; and Nie, F. 2018. New 12, 1-norm Relaxation of Multi-Way Graph Cut for Clustering. In *AAAI Conference on Artificial Intelligence*, 1–8.
- Yang, X.; Deng, C.; Wei, K.; Yan, J.; and Liu, W. 2020. Adversarial Learning for Robust Deep Clustering. In *Advances in Neural Information Processing Systems*, 9098–9108.
- Ye, M.; and Guo, Y. 2019. Progressive Ensemble Networks for Zero-Shot Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11728–11736.
- Zhu, P.; Wang, H.; and Saligrama, V. 2019. Generalized Zero-Shot Recognition based on Visually Semantic Embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2995–3003.
- Zhu, Y.; Xie, J.; Tang, Z.; Peng, X.; and Elgammal, A. 2019. Semantic-Guided Multi-Attention Localization for Zero-Shot Learning. In *Advances in Neural Information Processing Systems*, 1–11.