

Domain-aware Visual Bias Eliminating for Generalized Zero-Shot Learning

Shaobo Min¹, Hantao Yao², Hongtao Xie^{1*}, Chaoqun Wang¹, Zheng-Jun Zha¹, and Yongdong Zhang¹
¹University of Science and Technology of China

²National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences
 {mbobo, cq14}@mail.ustc.edu.cn, hantao.yao@nlpr.ia.ac.cn, {htxie, zhazj, zhyd73}@ustc.edu.cn

Abstract

Generalized zero-shot learning aims to recognize images from seen and unseen domains. Recent methods focus on learning a unified semantic-aligned visual representation to transfer knowledge between two domains, while ignoring the effect of semantic-free visual representation in alleviating the biased recognition problem. In this paper, we propose a novel Domain-aware Visual Bias Eliminating (DVBE) network that constructs two complementary visual representations, i.e., semantic-free and semantic-aligned, to treat seen and unseen domains separately. Specifically, we explore cross-attentive second-order visual statistics to compact the semantic-free representation, and design an adaptive margin Softmax to maximize inter-class divergences. Thus, the semantic-free representation becomes discriminative enough to not only predict seen class accurately but also filter out unseen images, i.e., domain detection, based on the predicted class entropy. For unseen images, we automatically search an optimal semantic-visual alignment architecture, rather than manual designs, to predict unseen classes. With accurate domain detection, the biased recognition problem towards the seen domain is significantly reduced. Experiments on five benchmarks for classification and segmentation show that DVBE outperforms existing methods by averaged 5.7% improvement.

1. Introduction

Deep learning has made great progress in many vision tasks, such as image classification [16, 28, 29, 53, 54, 60], object detection [22, 38], and semantic segmentation [6, 58, 59]. However, these remarkable technologies depend heavily on large-scale data of various classes, and it is infeasible to annotate them all. As a consequence, generalized zero-shot learning (GZSL) has attracted increasing attention, which can recognize images from either seen or unseen classes.

*Corresponding author.

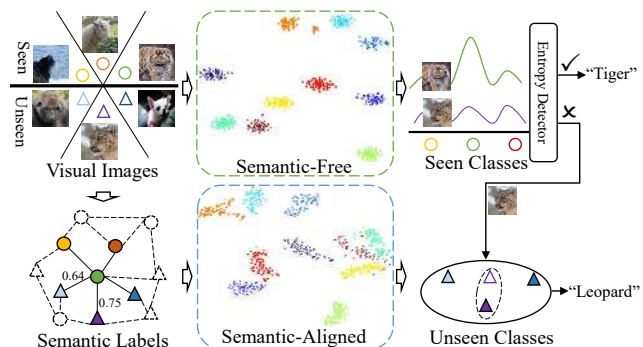


Figure 1. The inference framework of DVBE. Due to the weak discrimination of semantic-aligned visual representation, DVBE constructs an extra semantic-free visual representation to improve seen class prediction and filter out unseen images correctly.

A general paradigm of GZSL is to align the visual images and semantic labels, e.g., category attributes [12], in a joint semantic space, where the recognition becomes a nearest neighbor searching problem [45, 48, 61]. In the joint space, semantic embeddings and visual representations are treated as class anchors and input queries, respectively [1, 9]. However, the provided semantic labels are usually not discriminative [9, 37, 61], some of which are much similar, e.g., 0.75 cosine similarity between “Tiger” and “Leopard” in Figure 1. Thus, after semantic-visual alignment, the inter-class divergence between visual representations is inevitably weakened, which further results in the biased recognition problem [9, 10, 41], i.e., unseen images tend to be recognized as seen classes.

To address the above issue, recent methods [1, 5, 24, 41, 61, 62] focus on enlarging the visual representation discrimination. For example, SP-AEN [5] preserves the discriminative visual topology by reconstructing visual representations back to images. DLFZRL [41] factorizes an image representation into several latent components to improve model generalization. AREN [50] and VSE [61] explore part-based embeddings to capture detailed visual clues. Although being effective, the visual representations of these methods are ultimately aligned with semantic labels for

knowledge transfer. Therefore, they still suffer from degraded visual discrimination by semantic-visual alignment. Different from the previous methods, we consider learning an extra semantic-free visual representation in GZSL. Supervised with only class labels, the highly discriminative semantic-free representation can not only improve seen class prediction but also filter out unseen images based on seen class prediction entropy, *i.e.*, domain detection, as shown in Figure 1, though it cannot be transferred to the unseen domain. Once the domain detection is accurate, the filtered unseen images can be recognized in the domain-aware searching space, thereby alleviating the biased recognition problem.

Based on the above discussion, we propose a novel Domain-aware Visual Bias Eliminating (DVBE) network to construct two complementary visual representations, which are semantic-free and semantic-aligned, respectively. Explicitly, DVBE focuses on two main problems: a) how to boost the discrimination of semantic-free representation to improve domain detection accuracy; and b) how to design the semantic-visual alignment architecture for robust knowledge transfer. To this end, DVBE consists of two sub-modules, *i.e.*, adaptive margin second-order embedding (AMSE), and auto-searched semantic-visual embedding (autoS2V), as shown in Figure 2.

With only seen domain training data, AMSE targets to learn a semantic-free representation that is discriminative enough to classify seen images and filter out unseen images correctly. To this end, AMSE explores second-order visual statistics to capture subtle inter-class differences. To reduce the trivial information in high-dimensional second-order representations, a cross-attentive channel interaction is designed along both spatial and channel dimensions. Besides, adaptive margin constraints are imposed between decision boundaries of different classes to maximize the inter-class discrepancy. Consequently, AMSE significantly reduces the entropy of predicted seen class scores for seen domain images. Thus, unseen images of relatively high entropy can be correctly filtered out via an entropy-based detector. For the filtered unseen images, autoS2V is proposed to automatically search an optimal semantic-visual architecture via continuous architecture relaxation [30], which can produce robust semantic-aligned representation. Compared with manual designs [1, 61], autoS2V is more flexible in bridging the semantic-visual gap.

Experiments on object classification (averaged 4.3% improvement on CUB [44], AWA2 [46], aPY [12], and SUN [36]) and semantic segmentation (11% improvement on Pascal VOC [11]) prove the effectiveness of DVBE. Our contributions can be summarized by:

- We propose a novel Domain-aware Visual Bias Eliminating network that constructs two complementary visual representations, *i.e.*, semantic-free and semantic-

aligned, to tackle respective seen and unseen domains for unbiased GZSL.

- An adaptive margin second-order embedding is developed to generate highly discriminative semantic-free representation, which can improve seen class prediction and filter out unseen images accurately.
- An auto-searched semantic-visual embedding is designed to automatically search an optimal semantic-visual architecture for robust knowledge transfer.

2. Related Works

A general paradigm of GZSL is to align image representations and extra semantic information in a joint embedding space [5, 18, 51, 61]. The commonly used semantic information is within manually annotated category attributes [12, 35], word2vec [14, 57], and text descriptions [23]. Under this paradigm, the biased recognition problem [9, 10, 31, 41], *i.e.*, unseen images tend to be recognized as seen classes, is a major obstacle. One of the main reasons is that the provided semantic labels are usually not discriminative [9, 25, 61], which make the semantic-aligned visual representations from two domains hard to distinguish. To enlarge the visual discrimination, previous methods [33, 46, 61] replace the embedding space, spanned by semantic labels, with the high dimensional visual space. This prevents a few semantic embeddings of the seen domain from being nearest neighbors for most of the image samples [1]. Later, the auto-encoder architecture is widely adopted [20, 41] to preserve discriminative latent information. For example, SP-AEN [5] reconstructs visual representations back to images to retain the visual topology. PSR-ZSL [1] defines explicit semantic relations, *i.e.*, dissimilar, similar, and identical, during semantic reconstruction, and DLFZRL [41] exploits auto-encoder to factorize an image into several latent features and fuse them into a generalized one. Furthermore, to improve the description ability of visual representations, part-based embeddings are used in [50, 61, 62]. By automatically localizing informative regions, they can capture a lot of important local clues to improve semantic-aligned representations, which achieve robust knowledge transfer.

Recently, generative methods [13, 21, 34, 40, 47, 49] obtain impressive results by training powerful GANs [15] in the seen domain and directly generate massive unseen visual features from the semantic labels. This allows them to train a fully supervised classifier for two domains, which is robust to the biased recognition problem. In this paper, our focus is embedding-based GZSL, which requires no extra synthetic training data and complex training processing.

Although being effective, a commonality of the above embedding-based methods is that their visual representations are finally aligned with semantic labels for knowl-

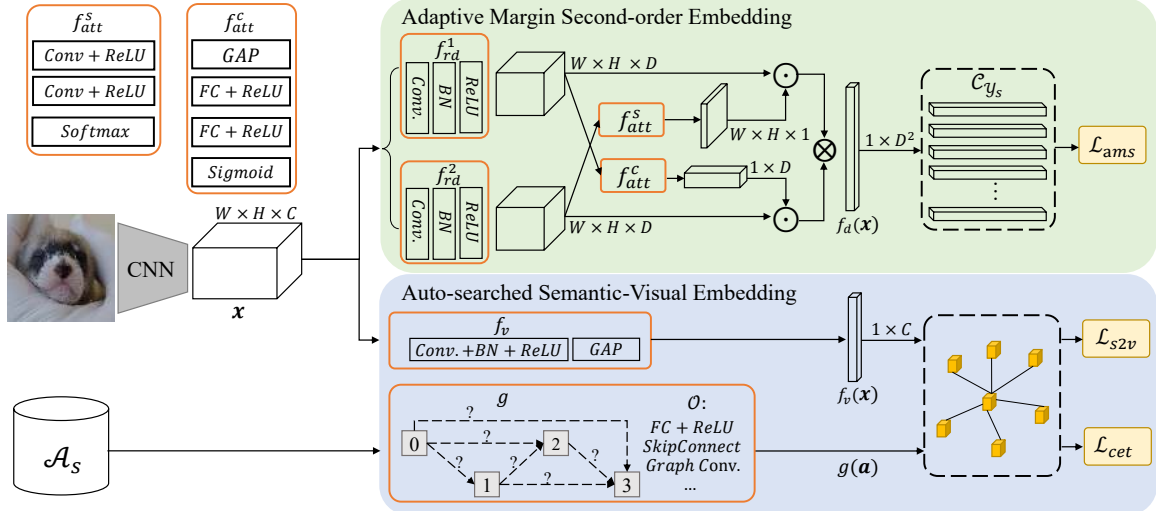


Figure 2. The training framework of DVBE with detailed implementation. *GAP* is global average pooling, and DVBE is trained with only seen domain data.

edge transfer, which are inevitably affected by the weak semantic discrimination. Differently, the proposed DVBE explores the totally semantic-free visual representation and find it can alleviate the biased recognition problem. Notably, COSMO [2] is the most similar work to us, which is also based on domain detection. However, COSMO focuses on designing an elaborate domain detector to improve domain detection. Differently, DVBE emphasizes on learning a discriminative semantic-free representation via AMSE to improve domain detection, and adopts a simple entropy-based detector for end-to-end training. The experiments show that our DVBE is superior to COSMO.

3. Domain-aware Visual Bias Eliminating

Generalized zero-shot learning aims to recognize images from seen and unseen domains. Formally, we define the seen domain data as $\mathcal{S} = \{\mathbf{x}_s, y_s, \mathbf{a}_s | \mathbf{x}_s \in \mathcal{X}_s, y_s \in \mathcal{Y}_s, \mathbf{a}_s \in \mathcal{A}_s\}$, where \mathbf{x}_s is the backbone feature for an image, y_s is the class label, and \mathbf{a}_s denotes the semantic label, such as class attributes or text descriptions. Similarly, the unseen domain data is defined as $\mathcal{U} = \{\mathbf{x}_u, y_u, \mathbf{a}_u | \mathbf{x}_u \in \mathcal{X}_u, y_u \in \mathcal{Y}_u, \mathbf{a}_u \in \mathcal{A}_u\}$, where $\mathcal{Y}_s \cap \mathcal{Y}_u = \emptyset$. Trained with seen domain data \mathcal{S} , GZSL targets to recognize a sample from either \mathcal{X}_s or \mathcal{X}_u .

A general paradigm of GZSL is to learn a semantic-visual space by minimizing:

$$\mathcal{L}_{s2v} = \sum_{\mathbf{x} \in \mathcal{X}_s} d(f_v(\mathbf{x}), g(\mathbf{a}_y)), \quad (1)$$

where y is the class label for sample \mathbf{x} , and \mathbf{a}_y is the corresponding semantic label for class y . $f_v(\cdot)$ and $g(\cdot)$ are the visual and semantic embedding functions to project \mathbf{x} and \mathbf{a}_y into the joint embedding space. The metric function

$d(\cdot, \cdot)$ measures the distance between $f_v(\mathbf{x})$ and $g(\mathbf{a}_y)$, e.g., cosine distance [1]. Based on an assumption that seen and unseen domains share the same semantic space, the semantic-visual alignment inferred from \mathcal{X}_s and \mathcal{A}_s can be directly transferred to \mathcal{X}_u and \mathcal{A}_u . Thus, the inference of GZSL becomes a nearest neighbor searching problem which can be defined by:

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

where $\mathbf{x} \in \mathcal{X}_s \cup \mathcal{X}_u$, and \hat{y} is the predicted class.

Although the semantic-visual alignment constructed by $f_v(\cdot)$ and $g(\cdot)$ can be adapted to \mathcal{U} , the visual discrimination of $f_v(\mathbf{x})$ is disturbed by the semantic labels. For example, if two categories have much similar semantic labels, the visual representations of images belonging to these two categories are also hard to distinguish, due to semantic-visual alignment in Eq. (1). Unfortunately, the provided semantic labels usually have a small inter-class difference between two domains [61], which makes the semantic-aligned $f_v(\mathbf{x}_s)$ and $f_v(\mathbf{x}_u)$ hard to distinguish and $f_v(\mathbf{x}_u)$ tend to be matched with \mathbf{a}_s .

3.1. Formulation

To tackle the above issues, we propose a novel Domain-aware Visual Bias Eliminating (DVBE) network. Besides the semantic-aligned visual representation $f_v(\mathbf{x})$ in Eq. (2), we introduce an extra semantic-free visual representation $f_d(\mathbf{x})$ that is trained on $\{\mathbf{x}_s, y_s\}$. Although $f_d(\mathbf{x})$ cannot be transferred to predict unseen classes without semantic relations, it can filter out the unseen images based on seen class prediction entropy, which enables $f_v(\mathbf{x})$ to focus on inferring the unseen classes. Based on the above discussion,

we converts Eq. (2) to:

$$\hat{y} = \begin{cases} \arg \max_{y \in \mathcal{Y}_s} \mathcal{C}_y(f_d(\mathbf{x})) & \text{if } \mathcal{H}(\mathcal{C}(f_d(\mathbf{x}))) \leq \tau, \\ \arg \min_{y \in \mathcal{Y}_u} d(f_v(\mathbf{x}), g(\mathbf{a}(y))) & \text{else,} \end{cases} \quad (3)$$

where $f_d(\mathbf{x})$ and $f_v(\mathbf{x})$ are semantic-free and semantic-aligned embedding functions, respectively. \mathcal{C} is a $|\mathcal{Y}_s|$ -way classifier. $\mathcal{H}(\cdot)$ measures the entropy of predicted scores of $f_d(\mathbf{x})$. Notably, a higher entropy denotes that the input image is more likely from the unseen domain, vice versa. In this paper, an image is regarded to be from the unseen domain once its entropy is higher than τ .

As only \mathcal{S} is available for training, we target to make $f_d(\mathbf{x})$ discriminative enough to produce a low entropy for most \mathbf{x}_s , thus many of the unseen images can be filtered out with relatively high entropy in Eq. (3). Notably, those correctly filtered unseen images \mathbf{x}_u are treated with domain-aware recognition in the searching space \mathcal{Y}_u by Eq. (3). Under this scheme, DVBE focuses on tackling two main problems: a) how to design $f_d(\cdot)$ that is highly discriminative for seen classes; and b) how to design $\{f_v(\cdot), g(\cdot)\}$ that can bridge the semantic-visual gap for correct unseen class prediction. Based on the above two problems, DVBE consists of two sub-modules, *i.e.*, Adaptive Margin Second-order Embedding and Auto-searched Semantic-Visual Embedding. Detailed architecture is shown in Figure 2.

3.2. Adaptive Margin Second-order Embedding

The previous works [2, 42] have shown that the out-of-distribution images can be filtered by applying an entropy-based detector once the visual representation is discriminative enough for the in-distribution classes. Therefore, we target to improve the visual discrimination of $f_d(\mathbf{x})$ to boost the domain detection in Eq. (3). To this end, an Adaptive Margin Second-order Embedding (AMSE) is developed by exploring discriminative second-order visual statistics and adaptive margin constraints.

Define the input feature as $\mathbf{x} \in R^{W \times H \times C}$, where W , H , and C are the width, height, and channels, respectively. For the convenience of bilinear operation, we reshape the feature $\mathbf{x} \in R^{W \times H \times C}$ to $\mathbf{x} \in R^{N \times C}$, where $N = W \times H$. Based on the feature \mathbf{x} , AMSE is proposed based on the well-known bilinear pooling [27]:

$$\mathbf{x}^{bp} = \mathbf{x} \otimes \mathbf{x}, \quad (4)$$

where \otimes denotes the local pairwise interaction along channel by $\mathbf{x}^{bp} = \sum_{n=1}^N \mathbf{x}_n^\top \mathbf{x}_n$, and bilinear description \mathbf{x}^{bp} need be reshaped to a feature vector for the classifier. However, besides the informative elements, the noisy elements in \mathbf{x} are also magnified by \otimes operator, *e.g.*, a noisy feature element will be interacted with all other elements.




Figure 3. A toy comparison between different Softmax strategies.

To reduce the negative effect of noisy features in Eq. (4), a cross-attentive channel interaction is designed to generate a compact representation $f_d(\mathbf{x})$ by:

$$f_d(\mathbf{x}) = [f_{att}^s(\mathbf{x}_2) \odot \mathbf{x}_1] \otimes [f_{att}^c(\mathbf{x}_1) \odot \mathbf{x}_2], \quad (5)$$

where $\mathbf{x}_1 = f_{rd}^1(\mathbf{x})$ and $\mathbf{x}_2 = f_{rd}^2(\mathbf{x})$ are two reduction layers to project \mathbf{x} into different compressed sub-spaces $R^{N \times D}$ ($D \ll C$). $f_{att}^s(\cdot)$ and $f_{att}^c(\cdot)$ are proposed to generate spatial and channel attention maps, respectively. \odot is the Hadamard product for element-wise multiplication. As shown in Figure 2, $f_{att}^s(\cdot)$ and $f_{att}^c(\cdot)$ are deployed in a cross-attentive way to facilitate attention complementarity. Compared with Eq. (4), the cross-attentive channel interaction has two main advantages: a) two different attention functions can eliminate the trivial elements effectively; and b) the cross-attentive manner can improve inputs complementarity.

To further enlarge the discrimination of $f_d(\mathbf{x})$, AMSE then applies an adaptive margin Softmax to maximize the inter-class differences by:

$$\mathcal{L}_{ams} = - \sum_{\mathbf{x} \in \mathcal{X}_s} \log \frac{e^{\lambda W_y f_d(\mathbf{x})}}{e^{\lambda W_y f_d(\mathbf{x})} + \sum_{j \in \mathcal{Y}_s, j \neq y} e^{W_j f_d(\mathbf{x})}}, \quad (6)$$

where W is the classifier weight, and y is the groundtruth label. $\lambda \in (0, 1]$ is a margin coefficient. When $\lambda = 1$, \mathcal{L}_{ams} is the standard Softmax loss. Differently, \mathcal{L}_{ams} imposes a margin penalty λ between $f_d(\mathbf{x})$ is predicted as y and $\mathcal{Y} \setminus y$. For example, when $\lambda = 0.8$, the response between \mathbf{x} and W_y becomes $0.8W_y f_d(\mathbf{x})$, which puts a more strict constraint on the inter-class difference to compact samples.

In this work, λ is adaptively calculated according to sample difficulties by:

$$\lambda = e^{-(p_y(\mathbf{x})-1)^2/\sigma^2}, \quad (7)$$

where $p_y(\mathbf{x})$ is the predicted probability for \mathbf{x} belonging to the class y . Notably, Eq. (7) exactly follows a Gaussian distribution of mean value of 1 and variance of σ . When \mathbf{x} is an easy sample whose $p_y(\mathbf{x}) \approx 1$, \mathcal{L}_{ams} approximates the standard Softmax loss with $\lambda \approx 1$. Inversely, for the hard samples, λ will be small, which indicates a large margin constraint and more attention on this sample.




Figure 4. (a) The hand-designed architecture. (b) The learned task-specific architectures from autoS2V for CUB and AWA2, respectively. GC is the graph convolution.

Compared with fixed margin λ in [8], \mathcal{L}_{ams} adopts a non-linear λ in Eq. (7) to emphasize on those hard samples. This can prevent a model from being corrupted by a rigorous margin constraint, *e.g.*, applying a large margin to an easy sample that has been already well-recognized. A toy comparison under no margin, fixed margin, and adaptive margin Softmax is given in Figure 3. Different from the focal loss [26], \mathcal{L}_{ams} utilizes Gaussian margin constraints between decision boundaries, which is more interpretable. Finally, \mathcal{L}_{ams} enables $f_d(\mathbf{x})$ to have a larger inter-class difference than standard Softmax loss.

By exploring cross-attentive second-order statistics with adaptive margin Softmax, the semantic-free $f_d(\mathbf{x})$ learned by AMSE has strong visual discrimination for seen classes, which can improve seen class prediction and filter out the unseen images correctly in Eq. (3).

3.3. Auto-searched Semantic-Visual Embedding

For the unseen image recognition, we develop an auto-searched semantic-visual embedding (AutoS2V) to generate the semantic-aligned visual representation. Following the general semantic-visual paradigm, AutoS2V consists of $\{f_v(\cdot), g(\cdot)\}$ by minimizing \mathcal{L}_{s2v} in Eq. (1). In most existing methods [1, 5], $g(\cdot)$ is designed manually, *e.g.*, two FC+ReLU layers that is similar to $f_v(\cdot)$, as shown in Figure 4 (a). However, there exists an obvious gap between visual and semantic signals, *e.g.*, each element in \mathbf{a} corresponds to a specific class attribute. Thus, the processing between $g(\cdot)$ and $f_v(\cdot)$ should be different. In contrast to previous methods, AutoS2V automatically searches an optimal architecture for $g(\cdot)$ with task-specific operations, as shown in Figure 4 (b).

The $g(\cdot)$ is viewed as a directed acyclic graph with several nodes, as shown in Figure 2. The input node is omitted for simplification. Each node is connected with all previous nodes, and operations between two connected nodes are selected from an operation set \mathcal{O} . \mathcal{O} contains four type of operations: (1) fully connection; (2) graph convolution [19, 43, 52, 56]; (3) skip connection; and (4) none. Notably, graph convolution is a task-specific operation, which is an expert in modeling topology relationships among semantic

Table 1. The train and val sets contain seen classes, while the test set contains unseen classes.

Datasets	Attributes	$ \mathcal{Y}_s $	$ \mathcal{Y}_u $	Train	Val	Test
CUB [44]	312	150	50	7,057	1,764	2,967
AWA2 [46]	85	40	10	23,527	5,882	7,913
aPY [12]	64	20	12	5,932	1,483	7,924
SUN [36]	102	645	72	10,320	2,580	1,440

labels. To search the optimal operation $\mathcal{O}_{i,j}$ between node i and node j , we relax the choice selection as a softmax optimization problem, following [30] by:

$$\mathcal{O}_{i,j} = \arg \max_{c \in \mathcal{O}} \frac{\exp(\alpha_{i,j}^c)}{\sum_{o \in \mathcal{O}} \exp(\alpha_{i,j}^o)}, \quad (8)$$

where $\alpha_{i,j}^c$ is the score for operation c . $\alpha = \{\alpha_{i,j}^c\}$ contains all trainable architecture parameters for $g(\cdot)$. Compared with hand designs, the automatically searched $g(\cdot)$ is more flexible in bridging the semantic-visual gap, which improves the unseen class prediction.

3.4. Overall

Finally, the overall objective function for DVBE is:

$$\mathcal{L}_{all} = \mathcal{L}_{s2v} + \mathcal{L}_{ams} + \gamma \mathcal{L}_{cet}, \quad (9)$$

where $\mathcal{L}_{cet}(f_v(\mathbf{x}))$ is an auxiliary cross-entropy loss to avoid all $f_v(\mathbf{x})$ collapsing into a single point with a control weight γ . Besides the object classification task, we further extend the proposed DVBE to the challenging zero-shot semantic segmentation [3]. Detailed illustrations for segmentation and \mathcal{L}_{cet} are given in supplementary material.

4. Experiments

4.1. Experimental Settings

Datasets. For generalized zero-shot classification, four widely used benchmarks are evaluated, which are Caltech-USCD Birds-200-2011 (CUB) [44], SUN [36], Animals with Attributes2 (AWA2) [46], and Attribute Pascal and Yahoo (aPY) [12]. All these datasets use the provided category attributes as semantic labels. The split of seen/unseen classes follows the newly proposed settings in [46], which ensures that unseen classes are strictly unknown in the pre-trained models. The details about each dataset are listed in Table 1. For zero-shot semantic segmentation, Pascal VOC [11] is further evaluated, which follows the setting in [3] by splitting the 20 classes into 14/6 as seen and unseen classes. The semantic labels for VOC are generated by using ‘word2vec’ model [32], which is learned on Wikipedia corpus and can produce a 300 dimensional description. During training, the images that contain unseen class pixels are all removed.

Implementation Details. For classification, the backbone network is ResNet-101 [17], which is pre-trained on Im-

Table 2. Results of GZSL on four classification benchmarks. Generative methods (GEN) utilizes extra synthetic unseen domain data for training. Since many previous methods cannot be end-to-end trained, we define DVBE and DVBE* as fixing and finetuning the backbone weights, respectively. † indicates the prediction ensemble from global and local regions.

	Methods	CUB [44]			AWA2 [46]			aPY [12]			SUN [36]		
		MCA_u	MCA_s	H	MCA_u	MCA_s	H	MCA_u	MCA_s	H	MCA_u	MCA_s	H
GEN	FGN[47]	43.7	57.7	49.7	-	-	-	-	-	-	42.6	36.6	39.4
	SABR-I[37]	55.0	58.7	56.8	30.3	93.9	46.9	-	-	-	50.7	35.1	41.5
	f-VAEGAN-D2[49]	63.2	75.6	68.9	-	-	-	-	-	-	50.1	37.8	43.1
NON-GEN	CDL[18]	23.5	55.2	32.9	-	-	-	19.8	48.6	28.1	21.5	34.7	26.5
	PSR-ZSL[1]	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[5]	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[41]	-	-	37.1	-	-	45.1	-	-	31.0	-	-	24.6
	MLSE[9]	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[4]	55.8	52.3	53.0	48.5	83.2	61.3	-	-	-	47.9	30.4	36.8
	COSMO[2]	44.4	57.8	50.2	-	-	-	-	-	-	44.9	37.7	41.0
	PREN*[55]	32.5	55.8	43.1	32.4	88.6	47.4	-	-	-	35.4	27.2	30.8
	VSE-S*[61]	33.4	87.5	48.4	41.6	91.3	57.2	24.5	72.0	36.6	-	-	-
	AREN*†[50]	63.2	69.0	66.0	54.7	79.1	64.7	30.0	47.9	36.9	40.3	32.3	35.9
	DVBE	53.2	60.2	56.5	63.6	70.8	67.0	32.6	58.3	41.8	45.0	37.2	40.7
	DVBE*	64.4	73.2	68.5	62.7	77.5	69.4	37.9	55.9	45.2	44.1	41.6	42.8

ageNet [39]. The rest part of DVBE uses MSRA random initializer. For data augmentation, random cropping of 448×448 and horizontal flipping are used. The testing images are cropped from the center part, and averaged horizontal flipping results are reported. To train the auto-searched $g(\cdot)$, a two-stage training strategy is adopted. First, the backbone is fixed, and DVBE is trained by alternately updating α and the rest model weights. Then, α is fixed, and the whole DVBE is finetuned end-to-end. The batch size is 24, and the reduction channel is $D = 256$. The SGD optimizer is used with initial $lr = 0.001$, momentum=0.9, and 180 training epoch. The hyper-parameter of σ is set to be 0.5 for most cases, and τ will be analyzed later. For semantic segmentation, the backbone network is based on Deeplab v3+ [7], as well as the data augmentation and optimizer. Code is available at github.com/mboboGO/DVBE.

Metrics. Similar to [46], the harmonic mean, denoted by $H = (2MCA_u \times MCA_s) / (MCA_u + MCA_s)$, is used to evaluate the comprehensive performance of a model in two domains. MCA_s and MCA_u are the Mean Class Top-1 Accuracy for seen and unseen domains. For semantic segmentation, similar hIoU is defined following [3].

4.2. Comparison with state-of-the-art methods

Object Classification The results on CUB, AWA2, aPY, and SUN are given in Table 2. In terms of non-generative models, the proposed DVBE surpasses them by 2.5%, 4.7%, 8.3%, and 1.8% for CUB, AWA2, aPY, and SUN, respectively. Among the compared methods, AREN obtains the best performance on most of datasets by aggregating the results from global and local regions. Compared with AREN, DVBE only considers global representation and achieves an averaged 5.6% improvement on four datasets. Except for AREN, DVBE outperforms other global-based methods by an averaged 8.5% improvement. Notably, the

impressive performance of DVBE is mainly attributed to the large improvements in the unseen domain, *e.g.*, 8.0% and 7.9% improvements of MCA_u on AWA2 and aPY. This demonstrates the effectiveness of DVBE in alleviating the biased recognition problem.

Among these methods, COSMO [2] is the most similar work to DVBE with domain detection mechanism. From Table 2, we can see that the proposed DVBE outperforms COSMO by 6.3% improvement on CUB and comparable performance on SUN. After end-to-end training, DVBE surpasses COSMO by respective 18.3% and 1.8% improvements on CUB and SUN. The main reason is that AMSE can significantly improve the visual discrimination of semantic-free representation, thereby resulting in more accurate domain detection than COSMO. Besides, DVBE adopts a simple entropy-based domain detector, which makes DVBE end-to-end trainable. Notably, the relatively small improvement on SUN dataset is because that each class contains fewer images than other datasets, which makes it hard to obtain highly discriminative $f_d(x)$.

As the generative methods employ the prior unseen semantic labels to generate massive synthetic data, they can obtain better performance than the existing non-generative methods. In terms of brand-new f-VAEGAN-D2 [49], we can observe that DVBE still obtains comparable performance without using unseen domain knowledge.

Semantic Segmentation We further evaluate the proposed method on semantic segmentation task, and summarize the related results in Table 3. The results of SP-AEN are obtained by extending its public code to ZSL segmentation. ZS3Net [3] is a newly proposed generative-based method for zero-shot semantic segmentation, and we use their reported performance. Impressively, DVBE outperforms ZS3Net by a large margin of 11.0% in terms of hIoU. The reason is that the semantic labels from word2vec con-

Table 3. GZSL results on Pascal-VOC.

Methods	Seen		Unseen		Over	
	PA	mIoU	PA	mIoU	hPA	hIoU
Devise[14]	89.9	64.3	10.3	2.9	18.5	5.5
SP-AEN[5]	90.5	69.5	15.9	10.5	27.0	18.2
ZS3Net[3]	92.9	69.3	46.7	26.1	62.2	37.9
DVBE	89.7	53.1	60.2	45.4	72.0	48.9




Figure 5. Some results of zero-shot semantic segmentations. The second and third rows are SP-AEN and DVBE, respectively.

tain more noise than manually annotated labels in classification, which makes unseen data synthesis of ZS3Net much harder. Thus, ZS3Net suffers from serious biased recognition problem, *e.g.*, a large gap between 26.1% and 69.3% mIoU in unseen and seen domains. By learning discriminative semantic-free representation, the domain detection by DVBE can effectively alleviate biased recognitions. Especially for the unseen domain, the mIoU of DVBE is higher than ZS3Net, *i.e.*, 45.4% *vs* 26.1%. In addition, we visualize some results of DVBE and SP-AEN in Figure 5 to prove the effectiveness of semantic-free representations¹. It can be seen that, with accurate domain detection, the regions of unseen classes can be accurately localized by DVBE, and SP-AEN usually provides biased recognition towards the seen classes.

4.3. Ablation Studies


In this section, we analyze the effects of different components in DVBE. The BaseS2V is a semantic-visual baseline, which consists of only $\{f_v(\cdot), g(\cdot)\}$ based on Eq (2), and $g(\cdot)$ is implemented by two layer FC+ReLU. To evaluate the domain detection performance, we regard GZSL as a binary classification task and define R_s and R_u as the Recall in respective seen and unseen domains. For example, R_u represents the percentage of how many unseen images are recognized as unseen classes. Some results are reported in supplementary materials, including attention visualization and hyper-parameters evaluation.

Effects of τ . In this work, τ is a critical parameter to divide seen and unseen images based on the predicted seen class entropy in Eq. (3). We thus calculate the entropy statistics of seen and unseen domain images of DVBE, and give

¹We cannot visualize the results for ZS3Net due to unavailable codes.

Table 4. The effects of semantic-free representation. Base $f_d(\mathbf{x})$ indicates a Conv+ReLU layer. CSE denotes the cross-attentive second-order embedding. H_R is the harmonic mean of R_s and R_u . The backbone is fixed, and autoS2V is not adopted here.

Modules	MCA_s	MCA_u	H	R_s	R_u	H_R
BaseS2V	62.5	32.8	43.0	89.5	45.8	60.6
+ $f_d(\mathbf{x})$	51.9	46.6	49.1	61.4	72.8	66.6
+CSE	56.5	50.3	53.2	70.4	78.7	74.3
+ \mathcal{L}_{ams} (DVBE)	59.9	50.9	55.0	71.6	80.2	75.7



(a) The entropy distributions of seen and unseen images by DVBE.


(b) Evaluation of varying τ for MCA_s , MCA_u , and H .

Figure 6. The analysis of domain detection by DVBE.

some analyses about τ . From Figure 6 (a), we can see that the seen and unseen images from both CUB and AWA2 have separable entropy statistics. Thus, the entropy-based detector can well separate them apart using an appropriate threshold of τ . Usually, a higher τ means that DVBE may regard more unseen images as the seen domain, vice versa. Thus, different τ has different effects on domain detection. As shown in Figure 6 (b), when increasing τ , MCA_s and MCA_u have opposite changing tendencies. By taking H into consideration, we set τ to achieve a trade-off between MCA_u and MCA_s . Since different datasets have different category numbers and difficulties, τ is set to be 2.1, 2.1, 1.5, and 3.9 for respective CUB, AWA2, aPY, and SUN.

Effects of semantic-free visual representation. We then explore the effects of semantic-free visual representation in GZSL. Compared with BaseS2V, DVBE learns complementary semantic-free visual representation $f_d(\mathbf{x})$ to predict seen classes and filter out the unseen images. As shown in Table 4, by adding the semantic-free $f_d(\mathbf{x})$ to BaseS2V, H_R is improved from 60.6% to 66.6%. Notably, H_R measures the domain detection performance of a GZSL model, thus the improved H_R proves that the semantic-free $f_d(\mathbf{x})$ can effectively alleviate the biased recognition problem, and H is improved from 43.0% to 49.1%. When further improving the visual discrimination of $f_d(\mathbf{x})$ via CSE and \mathcal{L}_{ams} , both H_R and H are increasing. Besides, we also visualize the feature distributions before and after applying AMSE to

Table 5. Analysis of cross-attentive channel interaction on CUB.

	\otimes	$f_{att}^s(\cdot)$	$f_{att}^c(\cdot)$	cross	MCA_s	MCA_u	H
$f_d(\mathbf{x})$	✓				51.9	46.6	49.1
	✓	✓	✓		53.8	48.1	50.8
	✓	✓	✓	✓	56.5	50.3	53.2




Figure 7. Representation distributions with/without semantics on AWA2. The dashed circles indicate seen classes.

$f_d(\cdot)$ in Figure 7 (a) and (b). It can be seen that AMSE can significantly improve the visual discrimination of \mathbf{x} . Consequently, DVBE outperforms BaseS2V by 12.0% gain.

Cross-attentive channel interaction. We next evaluate the effects of different components in cross-attentive channel interaction. As shown in Table 5, using the bilinear pooling operation \otimes obtains 1.7% improvement of H upon the base $f_d(\mathbf{x})$. Further, by using $f_{att}^s(\cdot)$ and $f_{att}^c(\cdot)$ to generate attentive features, H is improved from 50.8% to 52.4%, which shows that the redundancy reduction can boost the visual discrimination of $f_d(\mathbf{x})$. Finally, the cross-attentive manner also brings improvement by enhancing the inputs complementarity for second-order representation.

Adaptive margin Softmax. We further analyze the effect of adaptive margin Softmax \mathcal{L}_{ams} . As shown in Table 6, the adaptive margin Softmax obtains a higher performance than standard and fixed margin Softmax. Specifically, imposing a fixed margin on Softmax obtains a lower performance than standard Softmax. Different from using a fixed margin, we adaptively adjust margin λ and achieve the best performance. In Figure 8, the training curves for different Softmax strategies show that \mathcal{L}_{ams} obtains both fast and stable convergence.

Auto-searched $g(\cdot)$. We finally demonstrate that using autoS2V can generate the optimal and task-related architectures $g(\cdot)$. As shown in Figure 9, the auto-searched $g(\cdot)$ brings further improvements in terms of MCA_u than hand-designed one, and the auto-searched $g(\cdot)$ for CUB and AWA2 have different structures in Figure 4. By further analyzing the obtained architectures, we have the following observations: a) the graph convolution, which is used to explore semantic topology, tends to be used once due to the over-smoothing problem [19]; and b) two branch embeddings are preferred, which is similar to the existing work [41]. After semantic alignment, the visual features from two domains are associated more tightly for knowl-

Table 6. Evaluations of different Softmax results on CUB.

Methods	MCA_s	MCA_u	H
Standard Softmax	56.5	50.3	53.2
Fixed Margin Softmax	58.2	47.1	52.1
Adaptive Margin Softmax	59.9	50.9	55.0




Figure 8. The training curves of different Softmax strategies.




Figure 9. The improvements obtained by autoS2V on four datasets.

edge transfer, as shown in Figure 7 (b) and (c).

5. Conclusion

In this paper, we propose a novel Domain-aware Visual Bias Eliminating (DVBE) network to solve the biased recognition problem in generalized zero-shot learning (GZSL). Different from previous methods that focus on semantic-aligned representations, we consider the effect of semantic-free representation in alleviating the biased recognition by detecting unseen images based on seen class prediction entropy. To further boost the visual discrimination of semantic-free representations, we develop an adaptive margin second-order embedding via cross-attentive channel interaction and adaptive margin Softmax. Besides, we automatically search the optimal semantic-visual architectures to produce robust semantic-aligned representation. Consequently, by exploring complementary visual representations, *i.e.*, semantic-free and semantic-aligned, DVBE outperforms existing methods by a large margin in terms of both classification and segmentation.

In the future, we target to explore elaborate domain detector to improve DVBE.

Acknowledgments

This work is supported by the National Key Research and Development Program of China (2017YFC0820600), the National Nature Science Foundation of China (61525206, U1936210, 61902399), the Youth Innovation Promotion Association Chinese Academy of Sciences (2017209), and the Fundamental Research Funds for the Central Universities under Grant WK2100100030.

References

- [1] Yashas Annadani and Soma Biswas. Preserving semantic relations for zero-shot learning. In *CVPR*, pages 7603–7612, 2018.
- [2] Yuval Atzmon and Gal Chechik. Adaptive confidence smoothing for generalized zero-shot learning. In *CVPR*, pages 11671–11680, 2019.
- [3] Maxime Bucher, Tuan-Hung VU, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. In *NeurIPS*, pages 468–479, 2019.
- [4] Yannick Le Cacheux, Herve Le Borgne, and Michel Crucianu. Modeling inter and intra-class relations in the triplet loss for zero-shot learning. In *ICCV*, pages 10333–10342, 2019.
- [5] Long Chen, Hanwang Zhang, Jun Xiao, Wei Liu, and Shih-Fu Chang. Zero-shot visual recognition using semantics-preserving adversarial embedding networks. In *CVPR*, pages 1043–1052, 2018.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4):834–848, 2017.
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, pages 4690–4699, 2019.
- [9] Zhengming Ding and Hongfu Liu. Marginalized latent semantic encoder for zero-shot learning. In *CVPR*, pages 6191–6199, 2019.
- [10] Mohamed Elhoseiny and Mohamed Elfeki. Creativity inspired zero-shot learning. In *ICCV*, pages 5784–5793, 2019.
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010.
- [12] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *CVPR*, pages 1778–1785. IEEE, 2009.
- [13] Rafael Felix, Vijay BG Kumar, Ian Reid, and Gustavo Carneiro. Multi-modal cycle-consistent generalized zero-shot learning. In *ECCV*, pages 21–37, 2018.
- [14] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Tomas Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NeurIPS*, pages 2121–2129, 2013.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, pages 2672–2680, 2014.
- [16] Bin Gu, Miao Xin, Zhouyuan Huo, and Heng Huang. Asynchronous doubly stochastic sparse kernel learning. In *AAAI*, 2018.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [18] Huajie Jiang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Learning class prototypes via structure alignment for zero-shot recognition. In *ECCV*, pages 118–134, 2018.
- [19] Michael Kampffmeyer, Yinbo Chen, Xiaodan Liang, Hao Wang, Yujia Zhang, and Eric P Xing. Rethinking knowledge graph propagation for zero-shot learning. In *CVPR*, pages 11487–11496, 2019.
- [20] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *CVPR*, pages 3174–3183, 2017.
- [21] Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, and Piyush Rai. Generalized zero-shot learning via synthesized examples. In *CVPR*, pages 4281–4289, 2018.
- [22] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, pages 734–750, 2018.
- [23] Jimmy Lei Ba, Kevin Swersky, Sanja Fidler, et al. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *ICCV*, pages 4247–4255, 2015.
- [24] Kai Li, Martin Renqiang Min, and Yun Fu. Rethinking zero-shot learning: A conditional visual classification perspective. In *ICCV*, pages 3583–3592, 2019.
- [25] Yan Li, Junge Zhang, Jianguo Zhang, and Kaiqi Huang. Discriminative learning of latent features for zero-shot recognition. In *CVPR*, pages 7463–7471, 2018.
- [26] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017.
- [27] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *ICCV*, pages 1449–1457, 2015.
- [28] An-An Liu, Yu-Ting Su, Wei-Zhi Nie, and Mohan Kankanhalli. Hierarchical clustering multi-task learning for joint human action grouping and recognition. *IEEE TPAMI*, 39(1):102–114, 2016.
- [29] C. Liu, H. Xie, Z. Zha, L. Yu, Z. Chen, and Y. Zhang. Bidirectional attention-recognition model for fine-grained object classification. *IEEE TMM*, pages 1–1, 2019.
- [30] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018.
- [31] Jiawei Liu, Zheng-Jun Zha, Di Chen, Richang Hong, and Meng Wang. Adaptive transfer network for cross-domain person re-identification. In *CVPR*, pages 7202–7211, 2019.
- [32] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, pages 3111–3119, 2013.
- [33] Shaobo Min, Hantao Yao, Hongtao Xie, Zheng-Jun Zha, and Yongdong Zhang. Domain-specific embedding network for zero-shot recognition. In *ACM Multimedia*, pages 2070–2078. ACM, 2019.
- [34] Ashish Mishra, Shiva Krishna Reddy, Anurag Mittal, and Hema A Murthy. A generative model for zero shot learning using conditional variational autoencoders. In *CVPR Workshops*, pages 2188–2196, 2018.

- [35] Pedro Morgado and Nuno Vasconcelos. Semantically consistent regularization for zero-shot recognition. In *CVPR*, volume 9, page 10, 2017.
- [36] Genevieve Patterson and James Hays. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*, pages 2751–2758. IEEE, 2012.
- [37] Akanksha Paul, Narayanan C Krishnan, and Prateek Munjal. Semantically aligned bias reducing zero shot learning. In *CVPR*, pages 7056–7065, 2019.
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015.
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Sathesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [40] Edgar Schonfeld, Sayna Ebrahimi, Samarth Sinha, Trevor Darrell, and Zeynep Akata. Generalized zero-and few-shot learning via aligned variational autoencoders. In *CVPR*, pages 8247–8255, 2019.
- [41] Bin Tong, Chao Wang, Martin Klinkigt, Yoshiyuki Kobayashi, and Yuuichi Nonaka. Hierarchical disentanglement of discriminative latent features for zero-shot learning. In *CVPR*, pages 11467–11476, 2019.
- [42] Apoorv Vyas, Nataraj Jammalamadaka, Xia Zhu, Dipankar Das, Bharat Kaul, and Theodore L Willke. Out-of-distribution detection using an ensemble of self supervised leave-out classifiers. In *ECCV*, pages 550–564, 2018.
- [43] Zhengjue Wang, Chaojie Wang, Hao Zhang, Zhibin Duan, Mingyuan Zhou, and Bo Chen. Learning dynamic hierarchical topic graph with graph convolutional network for document classification. In *AISTATS*, 2020.
- [44] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- [45] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *CVPR*, pages 69–77, 2016.
- [46] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE TPAMI*, 2018.
- [47] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *CVPR*, pages 5542–5551, 2018.
- [48] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning—the good, the bad and the ugly. In *CVPR*, pages 4582–4591, 2017.
- [49] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. f-vaegan-d2: A feature generating framework for any-shot learning. In *CVPR*, pages 10275–10284, 2019.
- [50] Guo-Sen Xie, Li Liu, Xiaobo Jin, Fan Zhu, Zheng Zhang, Jie Qin, Yazhou Yao, and Ling Shao. Attentive region embedding network for zero-shot learning. In *CVPR*, pages 9384–9393, 2019.
- [51] Ning Xu, An-An Liu, Yongkang Wong, Yongdong Zhang, Weizhi Nie, Yuting Su, and Mohan Kankanhalli. Dual-stream recurrent neural network for video captioning. *IEEE TCSVT*, 29(8):2482–2493, 2018.
- [52] Xu Yang, Cheng Deng, Feng Zheng, Junchi Yan, and Wei Liu. Deep spectral clustering using dual autoencoder network. In *CVPR*, pages 4066–4075, 2019.
- [53] Hantao Yao, Shiliang Zhang, Chenggang Yan, Yongdong Zhang, Jintao Li, and Qi Tian. Autabd: Automated bi-level description for scalable fine-grained visual categorization. *IEEE TIP*, 27(1):10–23, 2017.
- [54] Hantao Yao, Shiliang Zhang, Yongdong Zhang, Jintao Li, and Qi Tian. Coarse-to-fine description for fine-grained visual categorization. *IEEE TIP*, 25(10):4858–4872, 2016.
- [55] Meng Ye and Yuhong Guo. Progressive ensemble networks for zero-shot recognition. In *CVPR*, pages 11728–11736, 2019.
- [56] Hanwang Zhang, Zheng-Jun Zha, Yang Yang, Shuicheng Yan, and Tat-Seng Chua. Robust (semi) nonnegative graph embedding. *IEEE TIP*, 23(7):2996–3012, 2014.
- [57] Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *CVPR*, pages 2021–2030, 2017.
- [58] Rui Zhang, Sheng Tang, Yongdong Zhang, Jintao Li, and Shuicheng Yan. Scale-adaptive convolutions for scene parsing. In *ICCV*, pages 2031–2039, 2017.
- [59] Rui Zhang, Sheng Tang, Yongdong Zhang, Jintao Li, and Shuicheng Yan. Perspective-adaptive convolutions for scene parsing. *IEEE TPAMI*, 2019.
- [60] Feng Zheng, Xin Miao, and Heng Huang. Fast vehicle identification via ranked semantic sampling based embedding. In *IJCAI*, pages 3697–3703, 2018.
- [61] Pengkai Zhu, Hanxiao Wang, and Venkatesh Saligrama. Generalized zero-shot recognition based on visually semantic embedding. In *CVPR*, pages 2995–3003, 2019.
- [62] Yizhe Zhu, Jianwen Xie, Zhiqiang Tang, Xi Peng, and Ahmed Elgammal. Semantic-guided multi-attention localization for zero-shot learning. In *NeurIPS*, pages 14917–14927, 2019.