

TransZero: Attribute-Guided Transformer for Zero-Shot Learning

Shiming Chen^{1*}, Ziming Hong^{1*}, Yang Liu², Guo-Sen Xie³, Baigui Sun²,
Hao Li², Qinmu Peng¹, Ke Lu^{4,5}, Xinge You^{1†}

¹Huazhong University of Science and Technology ²Alibaba Group, Hangzhou, China

³Inception Institute of Artificial Intelligence ⁴University of Chinese Academy of Sciences ⁵Peng Cheng Laboratory, China

{shimingchen, pengqinmu, youxg}@hust.edu.cn {hoongzm, gsxiem}@gmail.com luk@ucas.ac.cn

Abstract

Zero-shot learning (ZSL) aims to recognize novel classes by transferring semantic knowledge from seen classes to unseen ones. Semantic knowledge is learned from attribute descriptions shared between different classes, which act as strong priors for localizing object attributes that represent discriminative region features, enabling significant visual-semantic interaction. Although some attention-based models have attempted to learn such region features in a single image, the transferability and discriminative attribute localization of visual features are typically neglected. In this paper, we propose an attribute-guided Transformer network, termed TransZero, to refine visual features and learn attribute localization for discriminative visual embedding representations in ZSL. Specifically, TransZero takes a feature augmentation encoder to alleviate the cross-dataset bias between ImageNet and ZSL benchmarks, and improves the transferability of visual features by reducing the entangled relative geometry relationships among region features. To learn locality-augmented visual features, TransZero employs a visual-semantic decoder to localize the image regions most relevant to each attribute in a given image, under the guidance of semantic attribute information. Then, the locality-augmented visual features and semantic vectors are used to conduct effective visual-semantic interaction in a visual-semantic embedding network. Extensive experiments show that TransZero achieves the new state of the art on three ZSL benchmarks. The codes are available at: <https://github.com/shiming-chen/TransZero>.

Introduction

Inspired by human cognitive competence, zero-shot learning (ZSL) was proposed to recognize new classes during learning by exploiting the intrinsic semantic relatedness between seen and unseen classes (Larochelle, Erhan, and Bengio 2008; Palatucci et al. 2009; Lampert, Nickisch, and Harmeling 2009). In ZSL, there are no training samples available for unseen classes in the test set, and the label spaces for the training set and test set are disjoint from each other. Thus, the key task for ZSL is to learn discriminative visual features for conducting effective visual-semantic interactions

*These authors contributed equally.

†Corresponding author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

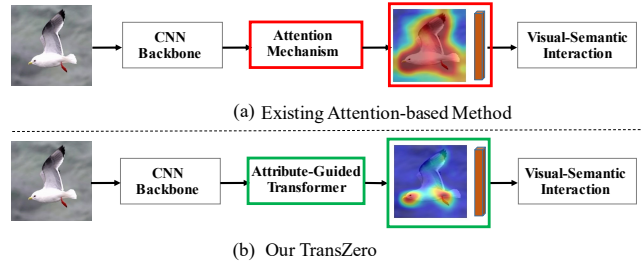


Figure 1: Motivation illustration. (a) Existing attention-based ZSL methods simply learn region embeddings (e.g., the whole bird body), neglecting the transferability and discriminative attribute localization (e.g., the distinctive bird body parts) of visual features; (b) Our TransZero reduces the entangled relationships among region features to improve their transferability and localizes the object attributes to represent discriminative region features, enabling significant visual-semantic interaction.

based on the semantic information (e.g., sentence embeddings (Reed et al. 2016), and attribute vectors (Lampert, Nickisch, and Harmeling 2014)), which are shared between the seen and unseen classes employed to support the knowledge transfer. According to their classification range, ZSL methods can be categorized into conventional ZSL (CZSL), which aims to predict unseen classes, and generalized ZSL (GZSL), which can predict both seen and unseen classes (Xian, Schiele, and Akata 2017).

To enable visual-semantic interactions, early ZSL methods attempt to build an embedding between seen classes and their class semantic vectors, and then classify unseen classes by nearest neighbor search in the embedding space. Since the embedding is only learned by seen class samples, these embedding-based methods usually overfit to seen classes under the GZSL setting (known as the bias problem). To tackle this problem, many generative ZSL methods have been proposed to generate samples of unseen classes by leveraging generative models (e.g., variational autoencoders (VAEs) (Arora et al. 2018; Schönfeld et al. 2019; Chen et al. 2021b), generative adversarial nets (GANs) (Xian et al. 2018, 2019; Chen et al. 2021a), and generative flows (Shen, Qin, and Huang 2020)) for data augmentation. Thus the ZSL task is

converted into a supervised classification problem.

Although these methods have achieved progressive improvement, they rely on global visual features which are insufficient for representing the fine-grained information of classes (e.g., red-legged of Kittiwake), since the discriminative information is contained in a few regions corresponding to a few attributes. Thus, the visual feature representations are limited, resulting in poor visual-semantic interactions. More recently, some attention-based models (Xie et al. 2019, 2020; Zhu et al. 2019; Xu et al. 2020; Yu et al. 2018; Liu et al. 2019) have attempted to learn more discriminative region features with the guidance of the semantic information, as shown in Fig. 1(a). However, these methods are limited in: i) They directly take the entangled region (grid) features for ZSL classification, which hinders the transferability of visual features from seen to unseen classes; ii) They simply learn region embeddings (e.g., the whole bird body), neglecting the importance of discriminative attribute localization (e.g., the distinctive bird body parts). Thus, properly improving the transferability and localizing the object attributes for enabling significant visual-semantic interaction in ZSL has become very necessary.

To tackle the above challenges, in this paper, we propose an attribute-guided Transformer, termed TransZero, which reduces the entangled relationships among region features to improve their transferability and localizes the object attributes to represent discriminative region features in ZSL, as shown in Fig. 1(b). Specifically, TransZero consists of an attributed-guided Transformer (AGT) that learns locality-augmented visual features and a visual-semantic embedding network (VSEN) that conducts visual-semantic interactions. In AGT, we first take a feature augmentation encoder to i) alleviate the cross-dataset bias between ImageNet and ZSL benchmarks, and ii) reduce the entangled relative geometry relationships between different regions for improving the transferability from seen to unseen classes. They are ignored by existing ZSL methods. To learn locality-augmented visual features, we employ a visual-semantic decoder in AGT to localize the image regions most relevant to each attribute in a given image, under the guidance of semantic attribute information. Then, the locality-augmented visual features and semantic vectors are used to enable visual-semantic interaction in VSEN. Extensive experiments show that TransZero achieves the new state of the art on three ZSL benchmarks. The qualitative results also demonstrate that TransZero refines visual features and provides attribute-level localization.

The main contributions of this paper are summarized as: **i)** We introduce a novel ZSL method, termed TransZero, which employs an attribute-guided Transformer to refine the visual features and learn the attribute localization for discriminative visual embedding representations. To the best of our knowledge, TransZero is the first work extending the Transformer to the ZSL task. **ii)** We propose a feature augmentation encoder to i) alleviate the cross-dataset bias between ImageNet and ZSL benchmarks, and ii) reduce the entangled relative geometry relationships between different regions to improve the transferability from seen to unseen classes. They are ignored by existing ZSL methods. **iii)** Extensive experiments demonstrate that TransZero achieves the new

state of the art on three ZSL benchmarks. We further qualitatively show that our TransZero refines the visual features and accurately localizes fine-grained parts for discriminative feature representations.

Related Work

Zero-Shot Learning. Early ZSL methods (Song et al. 2018; Li et al. 2018; Xian et al. 2018, 2019; Yu et al. 2020; Min et al. 2020; Han et al. 2021; Chen et al. 2021a; Chou, Lin, and Liu 2021; Han et al. 2021) focus on learning a mapping between the visual and semantic domains to transfer semantic knowledge from seen to unseen classes. They typically extract global visual features from pre-trained or end-to-end trainable networks. Typically, end-to-end models achieve better performance than pre-trained ones because they fine-tune the visual features, thus alleviating the cross-dataset bias between ImageNet and ZSL benchmarks (Chen et al. 2021a; Xian et al. 2019). However, these methods still usually yield relatively undesirable results, since they cannot efficiently capture the subtle differences between seen and unseen classes. More relevant to this work are the recent attention-based ZSL methods (Xie et al. 2019, 2020; Zhu et al. 2019; Xu et al. 2020; Liu et al. 2021) that utilize attribute descriptions as guidance to discover the more discriminative region (or part) features. Unfortunately, They simply learn region embeddings (e.g., the whole bird body) neglecting the importance of discriminative attribute localization (e.g., the distinctive bird body parts). Furthermore, the end-to-end attention models are also time-consuming when it comes to fine-tuning the CNN backbone. In contrast, we propose an attribute-guided Transformer to learn the attribute localization for discriminative region feature representations under non end-to-end ZSL model.

Transformer Model. Transformer models (Vaswani et al. 2017) have recently demonstrated excellent performance on a broad range of language and computer vision tasks, e.g., machine translation (Ott et al. 2018), image recognition (Dosovitskiy et al. 2021), video understanding (Gabeur et al. 2020), visual question answering (Zhang et al. 2021), etc. The success of Transformers can be mainly attributed to self-supervision and self-attention (Khan et al. 2021). The self-supervision allows complex models to be trained without the high cost of manual annotation, which in turn enables generalizable representations that encode useful relationships between the entities presented in a given dataset to be learned. The self-attention layers take the broad context of a given sequence into account by learning the relationships between the elements in the token set (e.g., words in language or patches in an image). Some methods (Gabeur et al. 2020; Cornia et al. 2020; Huang et al. 2019; Pan et al. 2020) have also shown that the transformer architecture can better capture the relationship between visual features and process sequences in parallel during training. Motivated by these, we design an attribute-guided Transformer that reduces the relationships among different regions to improve the transferability of visual features and learns the attribute localization for representing discriminative region features.

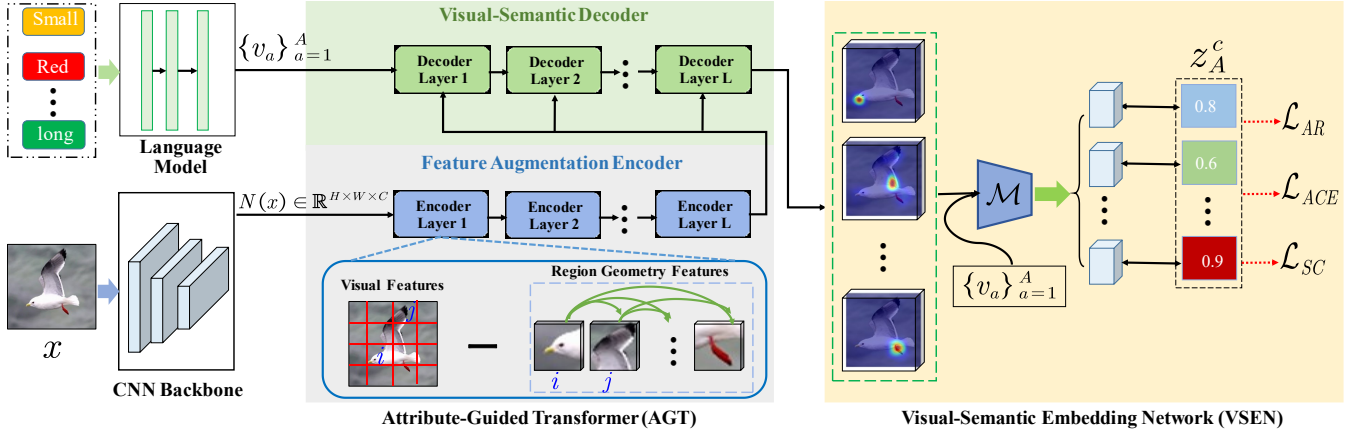


Figure 2: The architecture of the proposed TransZero model. TransZero consists of an attribute-guided Transformer (AGT) and a visual-semantic embedding network (VSEN). AGT includes a feature augmentation encoder that alleviates the cross-dataset bias between ImageNet and ZSL benchmarks and reduces the entangled geometry relationships between different regions for improving the transferability from seen to unseen classes, and a visual-semantic decoder that learns locality-augmented visual features based on the semantic attribute information. VSEN is used to enable significant visual-semantic interaction.

Proposed Method

We first introduce some notations and the problem definition. Assume that we have $\mathcal{D}^s = \{(x_i^s, y_i^s)\}$ as training data with C^s seen classes, where $x_i^s \in \mathcal{X}$ denotes the image i , and $y_i^s \in \mathcal{Y}^s$ is the corresponding class label. Another set of unseen classes C^u has unlabeled samples $\mathcal{D}^u = \{(x_i^u, y_i^u)\}$, where $x_i^u \in \mathcal{X}$ are the unseen class images, and $y_i^u \in \mathcal{Y}^u$ are the corresponding labels. A set of class semantic vectors of the class $c \in C^s \cup C^u = \mathcal{C}$ with A attributes $z^c = [z_1^c, \dots, z_A^c]^\top = \phi(y)$ helps knowledge transfer from seen to unseen classes. Note that we also use the semantic attribute vectors of each attribute $\mathcal{V}_A = \{v_a\}_{a=1}^A$ learned by a language model (i.e., GloVe (Pennington, Socher, and Manning 2014)) according to each word in attribute names. ZSL aims to predict the class labels $y^u \in \mathcal{Y}^u$ and $y \in \mathcal{Y} = \mathcal{Y}^s \cup \mathcal{Y}^u$ in the CZSL and GZSL settings, respectively, where $\mathcal{Y}^s \cap \mathcal{Y}^u = \emptyset$.

In this paper, we propose an attribute-guided Transformer network (termed TransZero) to refine the visual features and localize the object attributes for representing the discriminative region features under a non end-to-end model. This enables significant visual-semantic interaction in ZSL. As illustrated in Fig. 2, our TransZero includes an attribute-guided Transformer (AGT) and visual-semantic embedding network (VSEN). AGT refines the visual feature using a feature augmentation encoder, and learns locality-augmented visual features using a visual-semantic decoder. VSEN enables visual-semantic interaction for ZSL classification.

Attribute-Guided Transformer

Feature Augmentation Encoder. Since there is a cross-dataset bias between ImageNet and ZSL benchmarks (Chen et al. 2021a), we introduce a feature augmentation encoder to refine the visual features of ZSL benchmarks. Additionally, previous ZSL methods typically flatten the grid features (extracted from a CNN backbone) into a feature vector,

which is further used for generative models or embedding learning. However, such a feature vector implicitly entangles the feature representations among various regions in an image, which hinders their transferability from one domain to other domains (e.g., from seen to unseen classes) (Xu et al. 2020; Atzmon et al. 2020; Chen et al. 2021c). As such, we propose a feature-augmented scaled dot-product attention to further enhance the encoder layer by reducing the relative geometry relationships among the grid features.

To learn relative geometry features (Herdade et al. 2019; Zhang et al. 2021), we first calculate the relative center coordinates $(v_i^{\text{cen}}, t_i^{\text{cen}})$ based on the pair of 2D relative positions of the i -th grid $\{(v_i^{\text{min}}, t_i^{\text{min}}), (v_i^{\text{max}}, t_i^{\text{max}})\}$:

$$(v_i^{\text{cen}}, t_i^{\text{cen}}) = \left(\frac{v_i^{\text{min}} + v_i^{\text{max}}}{2}, \frac{t_i^{\text{min}} + t_i^{\text{max}}}{2} \right), \quad (1)$$

$$w_i = (v_i^{\text{max}} - v_i^{\text{min}}) + 1, \quad (2)$$

$$h_i = (t_i^{\text{max}} - t_i^{\text{min}}) + 1, \quad (3)$$

where $(v_i^{\text{min}}, t_i^{\text{min}})$ and $(v_i^{\text{max}}, t_i^{\text{max}})$ are the relative position coordinates of the top left corner and bottom right corner of the grid i , respectively.

Then, we construct region geometry features G_{ij} between grid i and grid j :

$$G_{ij} = \text{ReLU}(w_g^T g_{ij}), \quad (4)$$

where

$$g_{ij} = FC(r_{ij}), \quad r_{ij} = \begin{pmatrix} \log \left(\frac{|v_i^{\text{cen}} - v_j^{\text{cen}}|}{w_i} \right) \\ \log \left(\frac{|t_i^{\text{cen}} - t_j^{\text{cen}}|}{h_i} \right) \end{pmatrix}, \quad (5)$$

where r_{ij} is the relative geometry relationship between grids i and j , FC is a fully connected layer followed by a $ReLU$ activation, and w_g^T is a set of learnable weight parameters.

Finally, we subtract the region geometry features from the visual features in the feature-augmented scaled dot-product attention to provide a more accurate attention map, formally defined as:

$$Q^e = UW_q^e, K^e = UW_k^e, V^e = UW_v^e, \quad (6)$$

$$Z_{aug} = \text{softmax} \left(\frac{Q^e K^{e\top}}{\sqrt{d^e}} - G \right) V^e, \quad (7)$$

$$U \leftarrow U + Z_{aug}, \quad (8)$$

where Q, K, V are the query, key and value matrices, W_q^e, W_k^e, W_v^e are the learnable matrices of weights, d^e is a scaling factor, and Z_{aug} is the augmented features. $U \in \mathbb{R}^{HW \times C}$ are the packed visual features, which are learned from the flattened features embedded by a fully connected layer followed by a ReLU and a Dropout layer.

Visual-Semantic Decoder. Following the standard Transformer (Vaswani et al. 2017), our visual-semantic decoder takes a multi-head self-attention layer and feed-forward network (FFN) to build the decoder layer. The decoding process continuously incorporates visual information under the guidance of semantic attribute features \mathcal{V}_A . Thus, our visual-semantic decoder can effectively localize the image regions most relevant to each attribute in a given image. The multi-head self-attention layer uses the outputs of the encoder U as keys (K_t^d) and values (V_t^d) and a set of learnable semantic embeddings \mathcal{V}_A as queries (Q_t^d). It is defined as:

$$Q_t^d = \mathcal{V}_A W_{qt}^d, K_t^d = U W_{kt}^d, V_t^d = U W_{vt}^d, \quad (9)$$

$$\text{head}_t = \text{softmax} \left(\frac{Q_t^d K_t^{d\top}}{\sqrt{d^d}} \right) V_t^d, \quad (10)$$

$$\hat{F} = \parallel_{t=1}^T (\text{head}_t) W_o, \quad (11)$$

where $W_{qt}^d, W_{kt}^d, W_{vt}^d$ are the learnable weights, d^d is a scaling factor, and \parallel is a concatenation operation. Then, an FFN with two linear transformations followed a ReLU activation in between is applied to the attended features \hat{F} :

$$F = \text{ReLU} \left(\hat{F} W_1 + b_1 \right) W_2 + b_2, \quad (12)$$

where W_1, W_2, b_1 and b_2 are the weights and biases of the linear layers respectively, and F are the locality-augmented visual features.

Visual-Semantic Embedding Network

After generating locality-augmented visual features, we further map them into the semantic embedding space. To encourage the mapping to be more accurate, we take the semantic attribute vectors $\mathcal{V}_A = \{v_a\}_{a=1}^A$ as support, based on a mapping function (\mathcal{M}). Specifically, \mathcal{M} matches the locality-augmented visual features F with the semantic attribute information v_A :

$$\psi(x_i) = \mathcal{M}(F) = \mathcal{V}_A^\top W F, \quad (13)$$

where W is an embedding matrix that embeds F into the semantic attribute space. In essence, $\psi(x_i)[a]$ is an attribute

score that represents the confidence of having the a -th attribute in the image x_i . Given a set of semantic attribute vectors $\mathcal{V}_A = \{v_a\}_{a=1}^A$, TransZero attains a mapped semantic embedding $\psi(x_i)$.

Model Optimization

To achieve effective optimization, we employ the attribute regression loss, attribute-based cross-entropy loss and self-calibration loss to train TransZero.

Attribute Regression Loss. To encourage VSEN to accurately map visual features into their corresponding semantic embeddings, we introduce an attribute regression loss to constrain TransZero. Here, we regard visual-semantic mapping as a regression problem and minimize the mean square error between the ground truth attribute z^c and the embedded attribute score $\psi(x_i^s)$ of a set of sample $\{x_i^s\}_{i=1}^{n_b}$:

$$\mathcal{L}_{AR} = -\frac{1}{n_b} \sum_{i=1}^{n_b} \|\psi(x_i^s) - z^c\|_2^2. \quad (14)$$

Attribute-Based Cross-Entropy Loss. Since the associated image embedding is projected near its class semantic vector z^c when an attribute is visually present in an image, we take the attribute-based cross-entropy loss \mathcal{L}_{ACE} to optimize the parameters of the TransZero model, i.e., the dot product between the visual embedding and each class semantic vector is calculated to produce class logits. This encourages the image to have the highest compatibility score with its corresponding class semantic vector. Given a batch of n_b training images $\{x_i^s\}_{i=1}^{n_b}$ with their corresponding class semantic vectors z^c , \mathcal{L}_{ACE} is defined as:

$$\mathcal{L}_{ACE} = -\frac{1}{n_b} \sum_{i=1}^{n_b} \log \frac{\exp(\psi(x_i^s) \times z^c)}{\sum_{\hat{c} \in \mathcal{C}^s} \exp(\psi(x_i^s) \times z^{\hat{c}})}. \quad (15)$$

Self-Calibration Loss. Since \mathcal{L}_{AR} and \mathcal{L}_{ACE} optimize the model on seen classes, TransZero inevitably overfits to these classes, as also observed in (Zhu et al. 2019; Huynh and Elhamifar 2020a; Xu et al. 2020). To tackle this challenge, we further introduce a self-calibration loss \mathcal{L}_{SC} to explicitly shift some of the prediction probabilities from seen to unseen classes. \mathcal{L}_{SC} is thus formulated as:

$$\mathcal{L}_{SC} = -\frac{1}{n_b} \sum_{i=1}^{n_b} \sum_{c'=1}^{C^u} \log \frac{\exp(\psi(x_i^s) \times z^{c'} + \mathbb{I}_{[c' \in \mathcal{C}^u]})}{\sum_{\hat{c} \in \mathcal{C}} \exp(\psi(x_i^s) \times z^{\hat{c}} + \mathbb{I}_{[\hat{c} \in \mathcal{C}^u]})}, \quad (16)$$

where $\mathbb{I}_{[c' \in \mathcal{C}^u]}$ is an indicator function (i.e., it is 1 when $c' \in \mathcal{C}^u$, otherwise -1). Intuitively, \mathcal{L}_{ACE} encourages non-zero probabilities to be assigned to the unseen classes during training, which allows TransZero to produce a (large) non-zero probability for the true unseen class when given test samples from unseen classes.

Finally, we formulate the overall loss function of TransZero:

$$\mathcal{L}_{total} = \mathcal{L}_{ACE} + \lambda_{AR} \mathcal{L}_{AR} + \lambda_{SC} \mathcal{L}_{SC}, \quad (17)$$

where λ_{AR} and λ_{SC} are the weights to control their corresponding loss terms.

Methods	CUB				SUN				AWA2			
	CZSL	GZSL			CZSL	GZSL			CZSL	GZSL		
	acc	U	S	H	acc	U	S	H	acc	U	S	H
End-to-End												
QFSL (Song et al. 2018)	58.8	33.3	48.1	39.4	56.2	30.9	18.5	23.1	63.5	52.1	72.8	60.7
LDF (Li et al. 2018)	67.5	26.4	81.6	39.9	–	–	–	–	65.5	9.8	87.4	17.6
SGMA* (Zhu et al. 2019)	71.0	36.7	71.3	48.5	–	–	–	–	68.8	37.6	87.1	52.5
AREN* (Xie et al. 2019)	71.8	38.9	78.7	52.1	60.6	19.0	38.8	25.5	67.9	15.6	92.9	26.7
LFGAA* (Liu et al. 2019)	67.6	36.2	80.9	50.0	61.5	18.5	40.0	25.3	68.1	27.0	93.4	41.9
APN* (Xu et al. 2020)	72.0	65.3	69.3	67.2	61.6	41.9	34.0	37.6	68.4	57.1	72.4	63.9
Non End-to-End												
Generative Methods												
f-CLSWGAN (Xian et al. 2018)	57.3	43.7	57.7	49.7	60.8	42.6	36.6	39.4	68.2	57.9	61.4	59.6
f-VAEGAN-D2 (Xian et al. 2019)	61.0	48.4	60.1	53.6	64.7	45.1	38.0	41.3	71.1	57.6	70.6	63.5
OCD-CVAE (Keshari, Singh, and Vatsa 2020)	–	44.8	59.9	51.3	–	44.8	42.9	43.8	–	59.5	73.4	65.7
E-PGN (Yu et al. 2020)	72.4	52.0	61.1	56.2	–	–	–	–	73.4	52.6	83.5	64.6
Composer (Huyh and Elhamifar 2020b)	69.4	56.4	63.8	59.9	62.6	55.1	22.0	31.4	71.5	62.1	77.3	68.8
GCM-CF (Yue et al. 2021)	–	61.0	59.7	60.3	–	47.9	37.8	42.2	–	60.4	75.1	67.0
FREE (Chen et al. 2021a)	–	55.7	59.9	57.7	–	47.4	37.2	41.7	–	60.4	75.4	67.1
HSVA (Chen et al. 2021b)	62.8	52.7	58.3	55.3	63.8	48.6	39.0	43.3	–	59.3	76.6	66.8
Non-Generative Methods												
SP-AEN (Chen et al. 2018)	55.4	34.7	70.6	46.6	59.2	24.9	38.6	30.3	58.5	23.3	90.9	37.1
PQZSL (Li et al. 2019)	–	43.2	51.4	46.9	–	35.1	35.3	35.2	–	31.7	70.9	43.8
IIR (Cacheux, Borgne, and Crucianu 2019)	63.8	30.4	65.8	41.2	63.5	22.0	34.1	26.7	67.9	17.6	87.0	28.9
TCN (Jiang et al. 2019)	59.5	52.6	52.0	52.3	61.5	31.2	37.3	34.0	71.2	61.2	65.8	63.4
DVBE (Min et al. 2020)	–	53.2	60.2	56.5	–	45.0	37.2	40.7	–	63.6	70.8	67.0
DAZLE* (Huyh and Elhamifar 2020a)	66.0	56.7	59.6	58.1	59.4	52.3	24.3	33.2	67.9	60.3	75.7	67.1
TransZero (Ours)	76.8	69.3	68.3	68.8	65.6	52.6	33.4	40.8	70.1	61.3	82.3	70.2

Table 1: Results (%) of the state-of-the-art CZSL and GZSL modes on CUB, SUN and AWA2, including end-to-end and non end-to-end methods (generative and non-generative methods). The Symbol “–” indicates no results. The Symbol “*” denotes attention-based methods.

Zero-Shot Prediction

After training TransZero, we first obtain the embedding features of a test instance x_i in the semantic space i.e., $\psi(x_i)$. Then, we take an explicit calibration to predict the test label of x_i , which is formulated as:

$$c^* = \arg \max_{c \in \mathcal{C}^u / \mathcal{C}} \psi(x_i) \times z^c + \mathbb{I}_{[c \in \mathcal{C}^u]}. \quad (18)$$

Here, $\mathcal{C}^u / \mathcal{C}$ corresponds to the CZSL/GZSL setting respectively.

Experiments

Dataset Our extensive experiments are conducted on three popular ZSL benchmarks, including two fine-grained datasets (e.g., CUB (Welinder et al. 2010) and SUN (Patterson and Hays 2012)) and a coarse-grained dataset (e.g., AWA2 (Xian, Schiele, and Akata 2017)). CUB has 11,788 images of 200 bird classes (seen/unseen classes = 150/50) depicted with 312 attributes. SUN includes 14,340 images from 717 scene classes (seen/unseen classes = 645/72) depicted with 102 attributes. AWA2 consists of 37,322 images from 50 animal classes (seen/unseen classes = 40/10) depicted with 85 attributes.

Evaluation Protocols Following (Xian, Schiele, and Akata 2017), we measure the top-1 accuracy both in the CZSL and GZSL settings. In the CZSL setting, we simply predict the unseen classes to compute the accuracy of test samples, i.e., *acc*. In the GZSL setting, we compute the accuracy of the test samples from both the seen classes (denoted as \mathcal{S}) and unseen classes (denoted as \mathcal{U}). Meanwhile,

their harmonic mean (defined as $H = (2 \times S \times U) / (S + U)$) is also employed for evaluation in the GZSL setting.

Implementation Details We use the training splits proposed in (Xian et al. 2018). We take a ResNet101 pre-trained on ImageNet as the CNN backbone for extracting the feature map without fine-tuning. We use the SGD optimizer with hyperparameters (momentum = 0.9, weight decay = 0.0001) to optimize our model. The learning rate and batch size are set to 0.0001 and 50, respectively. We empirically set λ_{SC} to 0.3 and λ_{AR} to 0.005 for all datasets. The encoder and decoder layers are set to 1 with one attention head.

Comparison with State of the Art

Conventional Zero-Shot Learning. Here, we first compare our TransZero with the state-of-the-art methods in the CZSL setting. As shown in Table 1, our TransZero achieves the best accuracies of 76.8% and 65.6% on CUB and SUN, respectively. This shows that TransZero effectively learns the attribute-augmented region feature representations for distinguishing various fine-grained classes. As for the coarse-grained dataset, TransZero still obtains competitive performance, with a top-1 accuracy of 70.1%. Compared with other attention-based methods (e.g., SGMA (Zhu et al. 2019), AREN (Xie et al. 2019), APN (Xu et al. 2020)), TransZero obtains significant gains of over 4.8% and 4.0% on CUB and SUN, respectively. This demonstrates that the attribute localization representations learned by our TransZero are more discriminative than the region embeddings learned by the existing attention-based methods.

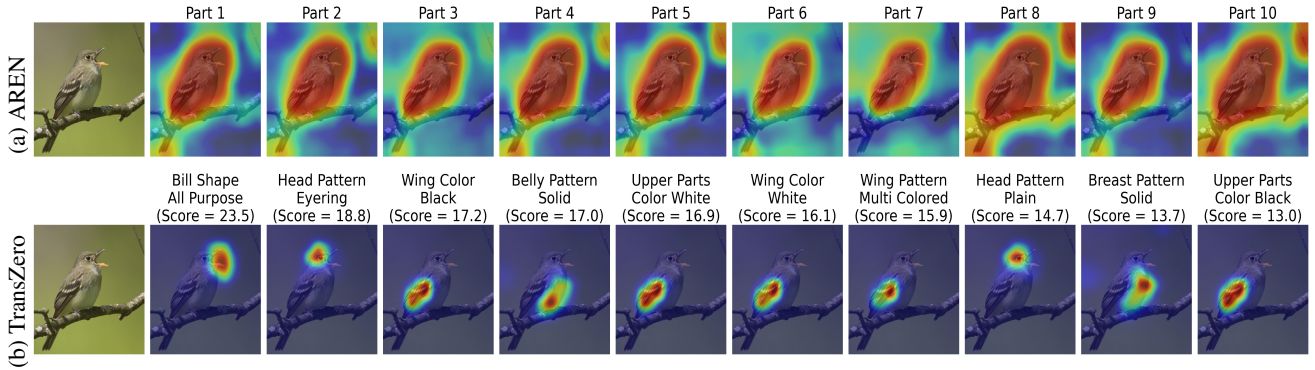


Figure 3: Visualization of attention maps for the attention-based method (i.e., AREN (Xie et al. 2019)) and our TransZero.

Method	CUB		SUN	
	acc	H	acc	H
TransZero w/o FAE	67.3	56.8	61.2	32.1
TransZero w/o FA	74.0	66.5	63.8	38.5
TransZero w/o DEC	62.3	53.7	58.3	31.6
TransZero w/o \mathcal{L}_{SC}	74.8	58.1	64.2	37.4
TransZero w/o \mathcal{L}_{AR}	74.5	67.3	64.1	39.1
TransZero (full)	76.8	68.8	65.6	40.8

Table 2: Ablation studies for different components of TransZero on the CUB and SUN datasets. “FAE” is the feature augmentation encoder, “FA” means feature augmentation, and “DEC” denotes visual-semantic decoder.

Generalized Zero-Shot Learning. Table 1 shows the results of different methods in the GZSL setting. The results show that the unseen accuracy (U) of all methods is usually lower than the seen accuracy (S) on the CUB and AWA2 datasets, i.e., $U < S$. Meanwhile, $U > S$ on the SUN dataset since the number of seen classes is much larger than the number of unseen classes.

We can see that most state-of-the-art methods achieve good results on seen classes but fail on unseen classes, while our method generalizes better to unseen classes with high unseen and seen accuracies. For example, TransZero achieves the best performance with harmonic mean of 68.8% and 70.2% on CUB and AWA2, respectively. We argue that the benefits of TransZero come from the fact that i) the feature augmentation encoder in AGT improves the discriminability and transferability of visual features, and ii) the self-calibration mechanism alleviates the bias problem. Finally, our TransZero also outperforms the attention-based methods by harmonic mean improvements of at least 1.6%, 3.2% and 3.1% on CUB, SUN and AWA2, respectively. This demonstrates the superiority and great potential of our attribute-guided Transformer for the ZSL task.

Ablation Study

To provide the further insight into TransZero, we conduct ablation studies to evaluate the effects of the feature augmentation encoder (denoted as FAE), feature augmenta-

tion in FAE (denoted as FA), visual-semantic decoder (denoted as DEC), self-calibration loss (i.e., \mathcal{L}_{SC}) and attribute regression loss (i.e., \mathcal{L}_{AR}). Our results are shown in Table 2. TransZero performs significantly worse than its full model when no feature augmentation encoder is used, i.e., the acc/harmonic mean drops by 9.5%/12.0% on CUB and 4.4%/8.7% on SUN. If we incorporate the encoder of the standard Transformer without feature augmentation, TransZero again achieves poor results compared to its full model, i.e., the acc/harmonic mean drops by 2.8%/2.3% and 1.8%/2.3% on CUB and SUN, respectively. When TransZero without visual-semantic decoder, its performances decreases dramatically on all datasets. Moreover, the self-calibration mechanism can effectively alleviate the bias problem, resulting in improvements in the harmonic mean of 10.7% and 3.4% on CUB and SUN, respectively. The attribute regression constraint further improves the performance of TransZero by directing VSEN to conduct effective visual-semantic mapping.

Qualitative Results

Visualization of Attention Maps. To intuitively show the effectiveness of our TransZero at learning locality-augmented visual features, we visualize the attention maps learned by the existing attention-based methods (e.g., AREN (Xie et al. 2019)) and TransZero. As shown in Fig. 3, AREN simply learns region embeddings for visual representations, e.g., the whole bird body, neglecting the fine-grained semantic attribute information. In contrast, our TransZero learns discriminative attribute localization for visual features by assigning high positive scores to key attributes (e.g., the ‘*bill shape all purpose*’ of the Acadian Flycatcher in Fig. 3). Thus, TransZero achieves significant performance both in seen and unseen classes.

t-SNE Visualizations. As shown in Fig. 4, we present the t-SNE visualization (Maaten and Hinton 2008) of visual features for (a) seen classes and (b) unseen classes on CUB, learned by the CNN backbone, TransZero encoder w/o FA, TransZero encoder, and TransZero decoder. When we incorporate the standard encoder into our TransZero, the visual features learned by the encoder are significantly improved compared to the original visual features extracted from the

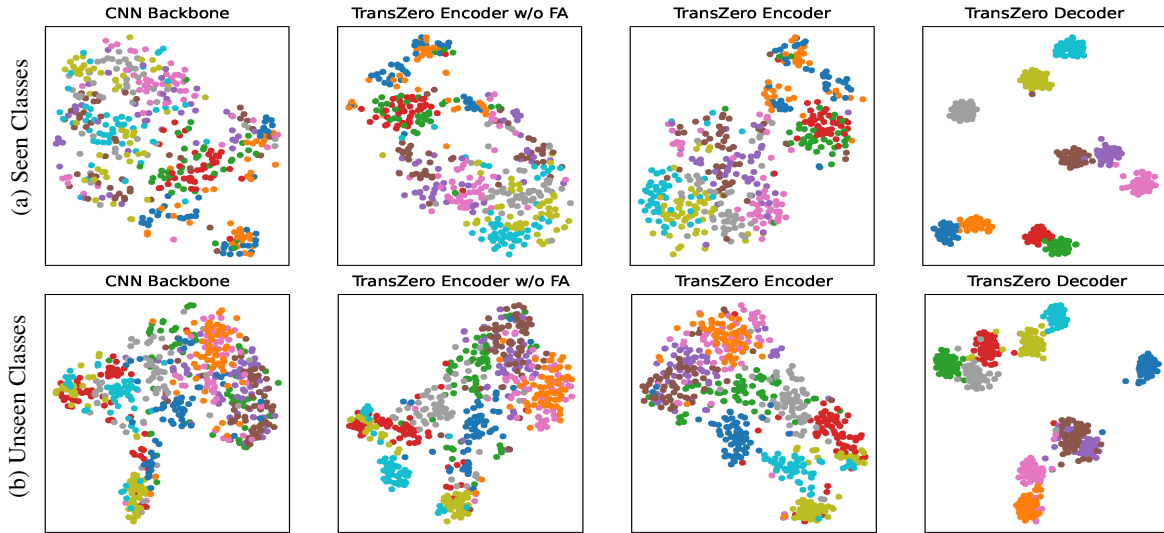


Figure 4: t-SNE visualizations of visual features for (a) seen classes and (b) unseen classes, learned by the CNN backbone, TransZero encoder w/o FA, TransZero encoder, and TransZero decoder. The 10 colors denote 10 different seen/unseen classes randomly selected from CUB.




Figure 5: The effects of λ_{AR} .

CNN Backbone (e.g., ResNet101). When we use the feature augmentation encoder to refine the original visual features, the quality of the unseen features is further enhanced. These results demonstrate that the encoder of TransZero effectively alleviates the cross-dataset bias problem and reduces the entangled relative geometry relationships among different regions, improving the transferability. Moreover, the visual-semantic decoder learns locality-augmented visual features for improving visual feature representations.

Hyperparameter Analysis

Effects of Loss Weight. λ_{AR} is employed to weigh the importance of the attribute regression loss, which directs the VSEN to conduct effective visual-semantic interaction. We try a wide range of λ_{AR} evaluated on CUB and SUN, i.e., $\lambda_{AR} = \{0.001, 0.005, 0.01, 0.05, 0.1, 0.5\}$. Results are shown in Fig. 5. When λ_{AR} is set to a large value, all evaluation protocols tend to drop. This is because the loss value of the attribute regression loss is too large, and thus the contributions of other losses are mitigated. When λ_{AR} is set to 0.005, our TransZero achieves the best performance.




Figure 6: The effects of λ_{SC} .

Effects of Loss Weight. λ_{SC} adjusts the weight of the self-calibration loss, which effectively alleviates the bias problem. As shown in Fig. 6, the accuracy on seen classes increases and the accuracy on unseen classes decreases when we increase λ_{SC} . Meanwhile, TransZero is insensitive to the self-calibration loss in the CZSL setting. We investigate a wide range of λ_{SC} on CUB and SUN to find an appropriate setting for λ_{SC} . Based on the results, we set λ_{SC} to 0.3 for all datasets.

Conclusion

In this paper, we propose a novel attribute-guided TransZero network for ZSL (termed TransZero). First, our TransZero employs a feature augmentation encoder to improve the discriminability and transferability of visual features by alleviating the cross-dataset problem and reducing the entangled region feature relationships. Meanwhile, a visual-semantic decoder is introduced to learn the attribute localization for locality-augmented visual feature representations. Secondly, a visual-semantic embedding network is used to enable effective visual-semantic interaction between the learned locality-augmented visual features and class semantic vectors. Extensive experiments on three popular benchmark datasets demonstrate the superiority of our approach.

We believe that our work also facilitates the development of other visual-and-language learning systems, e.g., natural language for visual reasoning.

Acknowledgements

This work is partially supported by NSFC (61772220), Special projects for technological innovation in Hubei Province (2018ACA135) and Key R&D Plan of Hubei Province (2020BAB027).

References

- Arora, G.; Verma, V.; Mishra, A.; and Rai, P. 2018. Generalized Zero-Shot Learning via Synthesized Examples. In *CVPR*, 4281–4289.
- Atzmon, Y.; Kreuk, F.; Shalit, U.; and Chechik, G. 2020. A causal view of compositional zero-shot recognition. In *NeurIPS*.
- Cacheux, Y. L.; Borgne, H.; and Crucianu, M. 2019. Modeling Inter and Intra-Class Relations in the Triplet Loss for Zero-Shot Learning. In *ICCV*, 10332–10341.
- Chen, L.; Zhang, H.; Xiao, J.; Liu, W.; and Chang, S. 2018. Zero-Shot Visual Recognition Using Semantics-Preserving Adversarial Embedding Networks. In *CVPR*, 1043–1052.
- Chen, S.; Wang, W.; Xia, B.; Peng, Q.; You, X.; Zheng, F.; and Shao, L. 2021a. FREE: Feature Refinement for Generalized Zero-Shot Learning. In *ICCV*.
- Chen, S.; Xie, G.-S.; Yang Liu, Y.; Peng, Q.; Sun, B.; Li, H.; You, X.; and Shao, L. 2021b. HSVA: Hierarchical Semantic-Visual Adaptation for Zero-Shot Learning. In *NeurIPS*.
- Chen, T.; Pu, T.; Xie, Y.; Wu, H.; Liu, L.; and Lin, L. 2021c. Cross-Domain Facial Expression Recognition: A Unified Evaluation Benchmark and Adversarial Graph Learning. *IEEE transactions on pattern analysis and machine intelligence*.
- Chou, Y.-Y.; Lin, H.-T.; and Liu, T.-L. 2021. Adaptive and Generative Zero-Shot Learning. In *ICLR*.
- Cornia, M.; Stefanini, M.; Baraldi, L.; and Cucchiara, R. 2020. Meshed-Memory Transformer for Image Captioning. In *CVPR*, 10575–10584.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Gabeur, V.; Sun, C.; Karteek, A.; and Schmid, C. 2020. Multi-modal Transformer for Video Retrieval. In *ECCV*.
- Han, Z.; Fu, Z.; Chen, S.; and Yang, J. 2021. Contrastive Embedding for Generalized Zero-Shot Learning. In *CVPR*.
- Herdade, S.; Kappeler, A.; Boakye, K.; and Soares, J. 2019. Image Captioning: Transforming Objects into Words. In *NeurIPS*.
- Huang, L.; Wang, W.; Chen, J.; and Wei, X.-Y. 2019. Attention on Attention for Image Captioning. In *ICCV*, 4633–4642.
- Huynh, D.; and Elhamifar, E. 2020a. Fine-Grained Generalized Zero-Shot Learning via Dense Attribute-Based Attention. In *CVPR*, 4482–4492.
- Huynh, D. T.; and Elhamifar, E. 2020b. Compositional Zero-Shot Learning via Fine-Grained Dense Feature Composition. In *NeurIPS*.
- Jiang, H.; Wang, R.; Shan, S.; and Chen, X. 2019. Transferable Contrastive Network for Generalized Zero-Shot Learning. In *ICCV*, 9764–9773.
- Keshari, R.; Singh, R.; and Vatsa, M. 2020. Generalized Zero-Shot Learning via Over-Complete Distribution. In *CVPR*, 13297–13305.
- Khan, S.; Naseer, M.; Hayat, M.; Zamir, S. W.; Khan, F.; and Shah, M. 2021. Transformers in Vision: A Survey. *arXiv preprint arXiv:2101.01169*.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2009. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 951–958.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2014. Attribute-Based Classification for Zero-Shot Visual Object Categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36: 453–465.
- Larochelle, H.; Erhan, D.; and Bengio, Y. 2008. Zero-data Learning of New Tasks. In *AAAI*, 646–651.
- Li, J.; Lan, X.; Liu, Y.; Wang, L.; and Zheng, N. 2019. Compressing Unknown Images With Product Quantizer for Efficient Zero-Shot Classification. In *CVPR*, 5458–5467.
- Li, Y.; Zhang, J.; Zhang, J.; and Huang, K. 2018. Discriminative Learning of Latent Features for Zero-Shot Recognition. In *CVPR*, 7463–7471.
- Liu, Y.; Guo, J.; Cai, D.; and He, X. 2019. Attribute Attention for Semantic Disambiguation in Zero-Shot Learning. In *ICCV*, 6697–6706.
- Liu, Y.; Zhou, L.; Bai, X.; Huang, Y.; Gu, L.; Zhou, J.; and Harada, T. 2021. Goal-Oriented Gaze Estimation for Zero-Shot Learning. In *CVPR*.
- Maaten, L. V. D.; and Hinton, G. E. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9: 2579–2605.
- Min, S.; Yao, H.; Xie, H.; Wang, C.; Zha, Z.; and Zhang, Y. 2020. Domain-Aware Visual Bias Eliminating for Generalized Zero-Shot Learning. In *CVPR*, 12661–12670.
- Ott, M.; Edunov, S.; Grangier, D.; and Auli, M. 2018. Scaling Neural Machine Translation. In *WMT*.
- Palatucci, M.; Pomerleau, D.; Hinton, G. E.; and Mitchell, T. M. 2009. Zero-shot Learning with Semantic Output Codes. In *NeurIPS*, 1410–1418.
- Pan, Y.; Yao, T.; Li, Y.; and Mei, T. 2020. X-Linear Attention Networks for Image Captioning. In *CVPR*, 10968–10977.
- Patterson, G.; and Hays, J. 2012. SUN attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*, 2751–2758.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global Vectors for Word Representation. In *EMNLP*.

- Reed, S.; Akata, Z.; Lee, H.; and Schiele, B. 2016. Learning Deep Representations of Fine-Grained Visual Descriptions. In *CVPR*, 49–58.
- Schönfeld, E.; Ebrahimi, S.; Sinha, S.; Darrell, T.; and Akata, Z. 2019. Generalized Zero- and Few-Shot Learning via Aligned Variational Autoencoders. In *CVPR*, 8239–8247.
- Shen, Y.; Qin, J.; and Huang, L. 2020. Invertible Zero-Shot Recognition Flows. In *ECCV*.
- Song, J.; Shen, C.; Yang, Y.; Liu, Y.; and Song, M. 2018. Transductive Unbiased Embedding for Zero-Shot Learning. *CVPR*, 1024–1033.
- Vaswani, A.; Shazeer, N. M.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NeurIPS*.
- Welinder, P.; Branson, S.; Mita, T.; Wah, C.; Schroff, F.; Belongie, S. J.; and Perona, P. 2010. Caltech-UCSD Birds 200. *Technical Report CNS-TR-2010-001, Caltech*.
- Xian, Y.; Lorenz, T.; Schiele, B.; and Akata, Z. 2018. Feature Generating Networks for Zero-Shot Learning. In *CVPR*, 5542–5551.
- Xian, Y.; Schiele, B.; and Akata, Z. 2017. Zero-Shot Learning — The Good, the Bad and the Ugly. In *CVPR*, 3077–3086.
- Xian, Y.; Sharma, S.; Schiele, B.; and Akata, Z. 2019. F-VAEGAN-D2: A Feature Generating Framework for Any-Shot Learning. In *CVPR*, 10267–10276.
- 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 *CVPR*, 9376–9385.
- Xie, G.-S.; Liu, L.; Jin, X.; Zhu, F.; Zhang, Z.; Yao, Y.; Qin, J.; and Shao, L. 2020. Region Graph Embedding Network for Zero-Shot Learning. In *ECCV*.
- Xu, W.; Xian, Y.; Wang, J.; Schiele, B.; and Akata, Z. 2020. Attribute Prototype Network for Zero-Shot Learning. In *NeurIPS*.
- Yu, Y.; Ji, Z.; Fu, Y.; Guo, J.; Pang, Y.; and Zhang, Z. 2018. Stacked Semantics-Guided Attention Model for Fine-Grained Zero-Shot Learning. In *NeurIPS*.
- Yu, Y.; Ji, Z.; Han, J.; and Zhang, Z. 2020. Episode-Based Prototype Generating Network for Zero-Shot Learning. In *CVPR*, 14032–14041.
- Yue, Z.; Wang, T.; Zhang, H.; Sun, Q.; and Hua, X. 2021. Counterfactual Zero-Shot and Open-Set Visual Recognition. In *CVPR*.
- Zhang, X.; Sun, X.; Luo, Y.; Ji, J.; Zhou, Y.; Wu, Y.; Huang, F.; and Ji, R. 2021. RSTNet: Captioning with Adaptive Attention on Visual and Non-Visual Words. In *CVPR*.
- Zhu, Y.; Xie, J.; Tang, Z.; Peng, X.; and Elgammal, A. 2019. Semantic-Guided Multi-Attention Localization for Zero-Shot Learning. In *NeurIPS*.