Goal-Oriented Gaze Estimation for Zero-Shot Learning

Yang Liu^{*1}, Lei Zhou^{*1}, Xiao Bai¹, Yifei Huang², Lin Gu^{3,2}, Jun Zhou⁴, Tatsuya Harada^{2,3}
¹Beihang University, ²The University of Tokyo, ³RIKEN AIP, Tokyo, Japan, ⁴Griffith University
{wickerboy, leizhou, baixiao}@buaa.edu.cn, hyf@iis.u-tokyo.ac.jp, lin.gu@riken.jp
jun.zhou@griffith.edu.au, harada@mi.t.u-tokyo.ac.jp

Abstract

Zero-shot learning (ZSL) aims to recognize novel classes by transferring semantic knowledge from seen classes to unseen classes. Since semantic knowledge is built on attributes shared between different classes, which are highly local, strong prior for localization of object attribute is beneficial for visual-semantic embedding. Interestingly, when recognizing unseen images, human would also automatically gaze at regions with certain semantic clue. Therefore, we introduce a novel goal-oriented gaze estimation module (GEM) to improve the discriminative attribute localization based on the class-level attributes for ZSL. We aim to predict the actual human gaze location to get the visual attention regions for recognizing a novel object guided by attribute description. Specifically, the task-dependent attention is learned with the goal-oriented GEM, and the global image features are simultaneously optimized with the regression of local attribute features. Experiments on three ZSL benchmarks, i.e., CUB, SUN and AWA2, show the superiority or competitiveness of our proposed method against the state-of-the-art ZSL methods. The ablation analysis on real gaze data CUB-VWSW also validates the benefits and accuracy of our gaze estimation module. This work implies the promising benefits of collecting human gaze dataset and automatic gaze estimation algorithms on high-level computer vision tasks. The code is available at <https://github.com/osierboy/GEM-ZSL>.

1. Introduction

With prior knowledge on seen classes, humans have a remarkable ability to recognize novel classes using shared and distinct attributes of both seen and unseen classes. Inspired by this cognitive competence, zero-shot learning (ZSL) was proposed as a challenging image classification setting to mimic the human cognitive process [31]. Given the semantic descriptions of both seen and unseen classes but only the

training images of seen classes, ZSL aims to classify test images of unseen classes.

Based on the classes that a model sees in the test phase, ZSL can be categorized into conventional or generalized setting. In conventional ZSL, the test images belong only to unseen classes. For the more practical and challenging generalized ZSL (GZSL) setting, the test images may belong to both seen and unseen classes. The semantic descriptions (attributes) are shared information between seen and unseen classes, which ensure the knowledge transferring. Early works [2, 3, 13, 39, 44] on ZSL build embedding between seen classes and their attributes. Then unseen classes are classified by the nearest neighbor search in the embedding space. These embedding based methods usually have a large bias towards seen classes under the GZSL setting, since the embedding is learned only by seen classes samples. To solve this problem, by leveraging the generative models [16, 29, 28], many feature generation approaches [30, 56, 61, 45, 46, 36] have been proposed to generate unseen classes, than convert ZSL into a conventional classification problem.

Most of the existing embedding or feature generation based methods extract global features from pre-trained or end-to-end trainable models. However, only the global image features cannot effectively represent the fine-grained information between seen and unseen images, which is important for ZSL. More recently, attention based end-to-end models [58, 67, 66, 59] have tried to exploit the semantic vector as guidance to learn more discriminative part features. However, they simply learn regions embedding of different attribute features but neglect the importance of discriminative attribute localization [49, 60].

Very interestingly, when facing an unseen object with the guidance of an attribute description, humans are capable of paying attention to parts of the object with discriminative attributes, which is the gaze behavior. Karessli *et al.* [26] proved that gaze data is useful in ZSL since they provide an effective prior that can naturally capture the localized discriminative attribute. Inspired by the human gaze mechanism, we propose a novel goal-oriented Gaze Estimation

^{*}Equal contribution.

Figure 1. Illustration of the proposed method. Our GEM-ZSL is an end-to-end trainable model with two main parts, a gaze estimation module and a plain ZSL backbone. For the gaze estimation module, the projected word vector of each attribute by **Word Encoder** is first utilized as query to guide the localized attribute learning. The ground truth attributes $\varphi(y)$ of specific seen class are simultaneously used as supervision in the **Attribute Localization** module. Then, an **Attention Transition** module is utilized to calibrate the attention regions by real gaze data with a designed gaze loss. Finally, the plain ZSL backbone learns the global image feature which is jointly optimized with the localized attribute, and the backbone accomplishes zero-shot recognition in a **Cosine Metric Space**.

Method for Zero-Shot Learning (GEM-ZSL). As shown in Figure 1, we first design an attribute description-oriented gaze estimation module (GEM) to learn different attribute regions. The GEM consists of three sub-modules, attention module (AM), attention transition (AT) module and attribute localization (AL) module. The AM is based on bilinear pooling which is widely used in the visual question answering task [14, 63]. We use the projected word vector of attributes as query to guide the learning process for the localization of the discriminative attribute. The ground truth attributes of specific seen class are simultaneously used as supervision by the mean squared error (MSE) of the AL module. Then an AT module is utilized to calibrate the attention regions by real gaze data (if available) with a designed gaze loss. In this way, we can learn the part feature with localized discriminative attribute that humans subconsciously pay attention to. Finally, the joint global features learned by an image encoder (IE), the local features and the class semantic embedding are used to learn a cosine metric space, which helps to reduce the intra-class variance and improve the recognition of unseen classes when compared to dot product similarity.

The contribution of this paper can be summarized as follows: 1. We propose a novel goal-oriented gaze estimation method to mimic the human cognitive process for recognizing unseen classes. With the guidance of attribute description, the proposed method can predict the human gaze that can be transformed to attribute attention for zero-shot recognition. 2. We demonstrate the effectiveness of our method for improving the localization of discriminative attributes,

which further enhances the discrimination of the global features for ZSL task. 3. Comprehensive experiments over three ZSL benchmarks, *i.e.*, CUB, AWA2 and SUN, show that our method can achieve superior or competitive performance compared with the state-of-the-art ZSL methods. In addition, the quantitative and qualitative results on gaze estimation experiment also validate the effectiveness of our GEM.

2. Related Works

2.1. Early ZSL

Early ZSL methods [2, 13, 44, 3, 54, 57, 64] focus on learning a mapping between visual and semantic spaces to transfer semantic knowledge from seen classes to unseen classes. However, these methods usually achieve relatively unsatisfactory results, since they adopt global features or exploit shallow models. In more recent, end-to-end deep models [39, 34, 47, 35] achieve better performance. These methods constrain loss on the attributes of seen classes to allow learning of more discriminative global features. However, they neglect to focus on the parts of features which are intrinsically discriminative for ZSL.

2.2. Part-based ZSL

More relevant to this work is the recently part-based ZSL methods [23, 58, 67, 59, 60] that utilize attribute descriptions as guidance to learn discriminative part features. These methods have achieved remarkable improvements on ZSL. In order to improve the GZSL, APN [60] applies cal-

ibrated stacking (CS) [8] in test phase to reduce the seen class scores by a constant factor, which can significantly improve the GZSL performance. RGEN [59] designs a balance loss by pursuing the maximum response consistency among seen and unseen outputs in the training phase, which requires the attributes of unseen classes during training. Compared with these methods, our GEM-ZSL improves the localization of discriminative attributes by goal-oriented gaze estimation for the first time. In addition, only the attributes of seen classes are used in the model training.

2.3. Gaze Estimation

Since human gaze directly represents attention and thus the important spatial regions, many works try to estimate where humans will look given an image. Earlier works [18, 1, 10] try to model saliency in the images that are likely to attract human attention. These methods typically use feature integration theory [50] to fuse various bottom-up cues. More recently researchers begin to use deep learning models for the estimation of human gaze [25, 40, 52]. However, these models only consider the bottom-up information which is not suitable for ZSL since they cannot focus on the fine-grained discriminative regions. Some newly proposed methods [21, 20] leverage top-down task-specific cues for estimating human gaze. However these methods are typically designed for some specific tasks and need large amount of training data. In this work, to solve the ZSL task, we exploit the attribute description as guidance to design a novel goal-oriented gaze estimation module.

3. The Proposed Method

In this section, we first define the problem setting, notations and then present the details of each module of our method.

3.1. Problem Setting and Notations

ZSL aims to recognize novel classes by transferring semantic knowledge from seen classes ($\mathcal{Y}^S$) to unseen classes ($\mathcal{Y}^U$). The image spaces of seen and unseen classes can be defined as $\mathcal{X} = \mathcal{X}^S \cup \mathcal{X}^U$. $\mathcal{S} = \{(\mathbf{x}, y, \varphi(y)) | \mathbf{x} \in \mathcal{X}^S, y \in \mathcal{Y}^S, \varphi(y) \in \phi^S\}$ denotes the training set, where $\mathbf{x}$ is an image in $\mathcal{X}^S$, y is its class label which is available during training, and $\varphi(y) \in \mathbb{R}^K$ is the class semantic embedding, *i.e.*, a class-level attribute vector annotated with K different visual attributes. The unseen testing set is $\mathcal{U} = \{(\mathbf{x}^u, u, \varphi(u)) | \mathbf{x}^u \in \mathcal{X}^U, u \in \mathcal{Y}^U, \varphi(u) \in \phi^U\}$, where u denote unseen class labels. The seen classes and unseen classes are disjoint, *i.e.*, $\mathcal{Y}^S \cap \mathcal{Y}^U = \emptyset$. Here, $\phi = \phi^S \cup \phi^U$ is used to transfer information between seen and unseen classes. In the conventional ZSL, the task is to predict the label of images from unseen classes, *i.e.*, $\mathcal{X}^U \rightarrow \mathcal{Y}^U$. However, in more realistic and challenging

setup of GZSL, the aim is to predict images from both seen and unseen classes, *i.e.*, $\mathcal{X} \rightarrow \mathcal{Y}^U \cup \mathcal{Y}^S$.

As shown in Figure 1, our end-to-end trained model consists of two parts, Gaze Estimation Module (GEM) and a plain ZSL backbone. Specifically, GEM consists of Word Encoder (WE) module, Attention Module (AM), Attribute Localization (AL) module and Attention Transition (AT) module. The plain ZSL backbone is an Image Encoder (IE) with a Cosine Metric Space for the nearest neighbor search. Firstly, IE is used to extract the global feature of the image. Then, GEM learns the part feature with localized discriminative attribute that humans subconsciously pay attention to. Finally, the joint global and local feature and the class semantic embedding are used to learn a cosine metric space.

3.2. Cosine Metric Learning

Visual feature extraction. We leverage the Image Encoder (IE) implemented by a convolutional neural network to map the seen class image $\mathbf{x}$ into a feature representation $f(\mathbf{x}) \in \mathbb{R}^{H \times W \times C}$, where H , W and C are the height, width and channel of the feature, respectively. Then, global average pooling is applied over the H and W to learn a global discriminative feature $h(\mathbf{x}) \in \mathbb{R}^C$.

Cosine similarity for classification. We use a linear layer $\mathbf{V} \in \mathbb{R}^{C \times K}$ to map the visual feature $h(\mathbf{x})$ into the semantic space. Different from previous work that [60] uses dot product to compute class logits of the projected visual feature and every class embedding, we consider the cosine similarity [15, 37] that can bound and reduce the variance of the neurons and thus result in models with better generalization capability [32]. The output of the cosine metric is scaled similarity score of the projected visual feature $h(\mathbf{x})$ and y -th class semantic embedding $\varphi(y)$. Then we define our classification score function as

$$p(y|\mathbf{x}) = \frac{\exp(\sigma \cos(h(\mathbf{x})^T \mathbf{V}, \varphi(y)))}{\sum_{\hat{y} \in \mathcal{Y}^S} \exp(\sigma \cos(h(\mathbf{x})^T \mathbf{V}, \varphi(\hat{y})))} \quad (1)$$

where σ is the scaling factor. The classification loss L_{CLS} is defined as

$$\mathcal{L}_{CLS} = -\log \frac{\exp(\sigma \cos(h(\mathbf{x})^T \mathbf{V}, \varphi(y)))}{\sum_{\hat{y} \in \mathcal{Y}^S} \exp(\sigma \cos(h(\mathbf{x})^T \mathbf{V}, \varphi(\hat{y})))} \quad (2)$$

We find empirically that using the image feature and class semantic embedding normalization in the cosine similarity helps to reduce the intra-class variance and improve the accuracy of unseen classes when compared to dot product similarity.

3.3. Gaze Estimation Module

Although cosine similarity shows better performance than the dot product similarity, the global features learned from cosine-based image encoder may still be biased to seen

Figure 2. The structure of attention module and attention transition module. It is based on the bilinear pooling mechanism. The feature map (key) and word vector (query) are shown as the shape of their tensor, *i.e.* $H \times W \times C$ for key and $K \times C$ for query. $\otimes$ denotes matrix multiplication. The softmax operation is performed on each row.

classes and weaken the discriminative attributes information. We propose a gaze estimation module (GEM) to learn the part feature with localized discriminative attribute that humans subconsciously pay attention to.

Word encoder. To learn local features, the attribute semantic vectors $\mathbf{e} = \{\mathbf{e}_k\}_{k=1}^K$ are utilized to guide the learning of the localized discriminative attribute, where $\mathbf{e}_k$ denotes the average GloVe [42] representations of words in the k -th attribute, *e.g.*, “brown eye” and “plain head”. Then, a single hidden layer MLP is used to convert the attribute word vector $\mathbf{e}$ into visual attribute feature $E(\mathbf{e}) \in \mathbb{R}^{K \times C}$.

Attention module. To acquire part features with localized discriminative attribute, we design an attribute descriptions-oriented attention module to learn the different attribute regions. As shown in Figure 2, AM is based on bilinear pooling mechanism which is widely used in the visual question answering task [14, 63]. We use the projected word vector of attributes $E(\mathbf{e})$ as a query to guide learning localized discriminative attribute. The inputs of AM are query $E(\mathbf{e}) \in \mathbb{R}^{K \times C}$ and key $f(\mathbf{x}) \in \mathbb{R}^{H \times W \times C}$. Firstly, AM reshapes the dimension of key into $HW \times C$ and transposes it into $C \times HW$. Then, AM performs matrix multiplication and softmax operation on query and key. The k -th attribute query captures related attribute part of image and produces $A(\mathbf{x})^k$ that contains the k -th localized attribute information. After transposing and reshaping the feature, the attribute feature map $A(\mathbf{x}) \in \mathbb{R}^{H \times W \times K}$ is acquired. The localized discriminative attribute is expected to concentrate to a peak region rather than disperses on other locations. We use Dis-

Figure 3. The predicted gaze map $g(\mathbf{x})$ and human gaze $G(\mathbf{x})$ matching method. We utilize the Hungarian algorithm which calculates the $l1$ distance between $g(\mathbf{x})$ and $G(\mathbf{x})$, and matches them one-to-one.

tance loss [65] on each attention map $A(\mathbf{x})^k$ to constrain the discriminative attribute localization, such that

$$\mathcal{L}_{Dis} = \sum_{k=1}^K \sum_{i=1}^H \sum_{j=1}^W A(\mathbf{x})_{i,j}^k (\|i - \tilde{i}\|^2 + \|j - \tilde{j}\|^2) \quad (3)$$

where $(\tilde{i}, \tilde{j}) = \arg \max_{i,j} A(\mathbf{x})^k$ denotes the coordinate for the maximum value in $A(\mathbf{x})^k$. We can learn a more concentrated attention map by this objective function.

Attribute localization. We predict the location of the attributes response value $a(\mathbf{x})$ by global max pooling over the H and W on $A(\mathbf{x})$. Then $a(\mathbf{x})$ is optimized by the Mean Square Error (MSE) with the supervision of the ground truth attributes $\varphi(y)$

$$\mathcal{L}_{MSE} = \|a(\mathbf{x}) - \varphi(y)\|_2^2 \quad (4)$$

where y is the ground truth class. We make the local features correspond to discriminative attributes by minimizing the MSE loss, which improves the discriminative attribute localization.

Attention transition. To acquire the part feature with discriminative attribute of human gaze, we propose an AT module to convert attribute feature map $A(\mathbf{x}) \in \mathbb{R}^{H \times W \times K}$ to gaze map $g(\mathbf{x}) \in \mathbb{R}^{H \times W \times D}$, where D represents the number of gaze heatmaps. We have true human gaze $G(\mathbf{x}) \in \mathbb{R}^{H \times W \times D}$ of image $\mathbf{x}$ to enable supervised learning. The AT is implemented by a learnable 1×1 convolution operation and a sigmoid function, which fuses different attributes information and converts them into a region of human attention. To optimize the gaze map that incorporates the discriminative attribute regions, we propose a new gaze loss.

Firstly, we use the Hungarian algorithm (Figure 3) to match $G(\mathbf{x})$ and $g(\mathbf{x})$, where the matched gaze map and its ground truth are termed $\tilde{g}(\mathbf{x})$ and $\tilde{G}(\mathbf{x})$, respectively. Then we use a binary cross-entropy loss across all the pixels as

$$\mathcal{L}_{Gaze} = -\frac{1}{D \times H \times W} \sum_{d=1}^D \sum_{i=1}^H \sum_{j=1}^W \{ \tilde{G}(\mathbf{x})_{i,j}^d \log(\tilde{g}(\mathbf{x})_{i,j}^d) + (1 - \tilde{G}(\mathbf{x})_{i,j}^d) \log(1 - \tilde{g}(\mathbf{x})_{i,j}^d) \} \quad (5)$$

Our full model optimizes the CNN backbone and the overall loss of the proposed model is defined as

$$\mathcal{L} = \mathcal{L}_{CLS} + \lambda_1 \mathcal{L}_{Dis} + \lambda_2 \mathcal{L}_{MSE} + \lambda_3 \mathcal{L}_{Gaze} \quad (6)$$

where λ_1 , λ_2 and λ_3 are hyper-parameters for Distance loss, Mean Square Error loss and Gaze loss, respectively. Particularly, $\lambda_3 = 0$ when gaze ground truth is not available. The joint training improves the representation of the discriminative attribute that is critical for zero-shot generalization.

3.4. Zero-Shot Recognition

After training of the full model, we use the learned cosine metric space for zero-shot recognition. The test image $\mathbf{x}$ is embedded to the cosine metric space by the visual-semantic embedding layer, and then the classifier searches for the class embedding $\varphi(\hat{u})$ with the highest compatibility via

$$\hat{u} = \arg \max_{u \in \mathcal{Y}^U} \cos(h(\mathbf{x})^T \mathbf{V}, \varphi(u)) \quad (7)$$

For the GZSL setting, the test images may belong to both seen and unseen classes. Since there are only seen classes during training phase, the predicted results of GZSL will have a large bias towards seen classes [8]. To mitigate this problem, we apply calibrated stacking (CS) [8] to reduce the seen class scores by a calibration factor γ . Specifically, the GZSL classifier is defined as

$$\hat{y} = \arg \max_{\tilde{y} \in \mathcal{Y}^U \cup \mathcal{Y}^S} (\sigma \cos(h(\mathbf{x})^T \mathbf{V}, \varphi(\tilde{y})) - \gamma \mathbb{I}[\tilde{y} \in \mathcal{Y}^S]) \quad (8)$$

where $\mathbb{I} = 1$ if $\tilde{y}$ is a seen class and 0 otherwise.

3.5. Implementation Details

The proposed GEM-ZSL is an end-to-end trainable model. The Image Encoder is ResNet101 [19] pretrained on ImageNet [11]. The SGD [6] optimizer is adopted in the model training. The momentum is set to 0.9, and the weight decay is 10^{-5} . The learning rate is 10^{-3} . We set λ_1 to 0.2 and λ_2 to 1.0. When we have gaze attention, λ_3 is set to 0.1. The factor γ is set to 3.5 for AWA2 and 0.7 for CUB and SUN. We use an episode-based training method to sample M categories and N images for each category in a mini-batch, we iterate 300 batches for each epoch, and train the model 20 epochs. We set $M = 16$ and $N = 2$ for all three datasets.

4. Experiments

We evaluated our framework on three widely used zero-shot learning benchmark datasets, including CUB-200-2011 (CUB) [53], SUN attribute (SUN) [41] and Animals with Attributes 2 (AWA2) [55]. The pre-defined attributes on each dataset were used as the semantic descriptors. Moreover, we adopted the Proposed Split (PS) [55] to divide all classes into seen and unseen classes on each dataset.

The performance of ZSL is evaluated by average per-class Top-1 (T1) accuracy. In GZSL, since the test set is composed of seen and unseen images, the Top-1 accuracy evaluated respectively on seen classes, denoted as $\mathbf{S}$, and unseen classes, denoted as $\mathbf{U}$. Their harmonic mean, defined as $\mathbf{H} = (2 \times \mathbf{S} \times \mathbf{U}) / (\mathbf{S} + \mathbf{U})$ [55], are used to evaluate the performance of GZSL.

4.1. Comparison with the State-of-the-Art

We selected recent state-of-the-art ZSL methods for comparison, which include methods without end-to-end training such as PSR [4], RN [48], SP-AEN [9], IIR [7], TCN [24], E-PGN [62], DA-ZLE [22], f-CLSWGAN [56], cycle-CLSWGAN [12], CADA-VAE [45], OCD-CVAE [27], RFF-GZSL [17], IZF [46], and LsrGAN [51], where the last seven methods are feature generation based models, and end-to-end methods QFSL [47], LDF [33], SGMA [67], AREN [58], LF-GAA [35], DVBE [38], RGEN [59], and APN [60].

Table 1 shows the results of different methods on three datasets. Our GEM-ZSL achieves competitive performance compared with the state-of-the-art methods. On CUB dataset, GEM-ZSL outperforms all the compared methods with a large margin for both ZSL and GZSL. Since CUB is a more challenging fine-grained dataset which requires local discriminative attributes, the results prove the effectiveness of our localized attribute learning model. For AWA2 dataset, our GEM-ZSL can also achieve competitive result which is only slightly lower than RGEN. However, the balance loss of RGEN requires the attributes of unseen classes during training which is not used in our GEM-ZSL. Without the balance loss, the harmonic mean for RGEN on AWA2 will dramatically decrease to 14.7%. On SUN dataset, the feature generation based model significantly outperforms the other methods. As SUN dataset contains more than 700 categories, the generative model can bring more features for generalization to unseen classes. Compared with the other non-generation based methods, the performance of our GEM-ZSL is competitive.

4.2. Ablation Study

Component analysis. We conducted ablation experiments to verify the effectiveness of the proposed modules. Table 2 shows the influence of each model compo-

Table 1. Results (%) of the state-of-the-art ZSL and GZSL. The first part is non end-to-end methods, the second part is feature generation methods and the third part is end-to-end methods. The best and the second best results are marked in red and blue, respectively.

Methods	CUB				SUN				AWA2			
	ZSL	GZSL			ZSL	GZSL			ZSL	GZSL		
	T1	U	S	H	T1	U	S	H	T1	U	S	H
PSR(CVPR'18) [4]	56.0	24.6	54.3	33.9	61.4	20.8	37.2	26.7	63.8	20.7	73.8	32.3
RN(CVPR'18) [48]	55.6	38.1	61.1	47.0	-	-	-	-	64.2	30.0	93.4	45.3
SP-AEN(CVPR'18) [9]	55.4	34.7	70.6	46.6	59.2	24.9	38.6	30.3	-	-	-	-
IIR(ICCV'19) [7]	63.8	55.8	52.3	53.0	63.5	47.9	30.4	36.8	67.9	48.5	83.2	61.3
TCN(ICCV'19) [24]	59.5	52.6	52.0	52.3	61.5	31.2	37.3	34.0	71.2	61.2	65.8	63.4
E-PGN(CVPR'20) [62]	72.4	52.0	61.1	56.2	-	-	-	-	73.4	52.6	83.5	64.6
DAZLE(CVPR'20) [22]	65.9	56.7	59.6	58.1	-	52.3	24.3	33.2	-	60.3	75.7	67.1
f-CLSWGAN(CVPR'18) [56]	57.3	43.7	57.7	49.7	60.8	42.6	36.6	39.4	-	-	-	-
cycle-CLSWGAN(ECCV'18) [12]	58.4	45.7	61.0	52.3	60.0	49.4	33.6	40.0	-	-	-	-
CADA-VAE(CVPR'19) [45]	-	51.6	53.5	52.4	-	47.2	35.7	40.6	-	55.8	75.0	63.9
OCD-CVAE(CVPR'20) [27]	60.3	44.8	59.9	51.3	63.5	44.8	42.9	43.8	71.3	59.5	73.4	65.7
RFF-GZSL(1-NN)(CVPR'20) [17]	-	50.6	79.1	61.7	-	56.6	42.8	48.7	-	-	-	-
IZF(ECCV'20) [46]	67.1	52.7	68.0	59.4	68.4	52.7	57.0	54.8	74.5	60.6	77.5	68.0
LsrGAN(ECCV'20) [51]	60.3	48.1	59.1	53.0	62.5	44.8	37.7	40.9	-	-	-	-
QFSL(CVPR'18) [47]	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(CVPR'18) [33]	67.5	26.4	81.6	39.9	-	-	-	-	-	-	-	-
SGMA(NeurIPS'19) [67]	71.0	36.7	71.3	48.5	-	-	-	-	-	-	-	-
AREN(CVPR'19) [58]	71.8	63.2	69.0	66.0	60.6	40.3	32.3	35.9	67.9	54.7	79.1	64.7
LFGAA(ICCV'19) [35]	67.6	36.2	80.9	50.0	61.5	18.5	40.0	25.3	68.1	27.0	93.4	41.9
DVBE(CVPR'20) [38]	-	64.4	73.2	68.5	-	44.1	41.6	42.8	-	62.7	77.5	69.4
RGEM(ECCV'20) [59]	76.1	60.0	73.5	66.1	63.8	44.0	31.7	36.8	73.6	67.1	76.5	71.5
APN(NeurIPS'20) [60]	72.0	65.3	69.3	67.2	61.6	41.9	34.0	37.6	68.4	56.5	78.0	65.5
GEM-ZSL(Ours)	77.8	64.8	77.1	70.4	62.8	38.1	35.7	36.9	67.3	64.8	77.5	70.6

Table 2. Results (%) of ZSL and GZSL ablation study on CUB, SUN and AWA2. The baseline is the Image Encoder and dot product distance with cross-entropy loss. We analyzed the performance of each module of our model.

Methods	CUB		SUN		AWA2	
	T1	H	T1	H	T1	H
Baseline	67.8	61.5	54.7	30.9	63.5	62.7
+ $\mathcal{L}_{MSE}$	69.9	63.1	56.1	31.5	64.3	64.9
+ $\mathcal{L}_{Dis}$	70.5	64.9	56.9	31.7	64.6	65.2
+ $\cos$	77.8	70.4	62.8	36.9	67.3	70.6

nent. Firstly, we train a baseline model that contains Image Encoder with cross-entropy loss. We use the dot product to compute class logits of the projected visual feature and every class embedding. Then, we add gaze estimation module, mean square error loss and distance loss functions gradually. Finally, the cosine similarity is added to the model. We can see that our proposed GEM improves the Top-1 accuracy (T1) of ZSL over the baseline consistently by 10.0% (CUB), 8.1% (SUN), 3.8% (AWA2), and the harmonic mean accuracy (H) of GZSL over the baseline

Table 3. Results (%) of ZSL ablation study for attention transition with gaze loss.

Methods	Set 1	Set 2	Set 3	Set 4	Avg.
GEM w/o $\mathcal{L}_{Gaze}$	43.5	45.5	42.1	40.7	42.9
GEM w $\mathcal{L}_{Gaze}$	44.9	45.2	43.6	41.2	43.7

by 8.9% (CUB), 6.0% (SUN), 7.9% (AWA2) respectively. Remarkably, when cosine distance is used to measure the similarity of the projected visual feature and class semantic embedding, the performance of our model can be greatly improved. Since the cosine metric can bound and reduce the variance of the neurons and thus result in models with better generalization capability [32].

To verify the influence of attention transition, we used CUB-VWSW gaze dataset with captured gaze points [26]. We performed Gaussian blur on the gaze points to obtain the gaze heatmap as gaze ground truth. Then we randomly selected two categories for each bird family as unseen categories. The other categories were used as seen categories. All seen categories were used to train the model with gaze

(a) **H**

(b) **T1**

Figure 4. The influence of σ with different values of ZSL (**T1**) and GZSL (**H**) results (%) on CUB, SUN and AWA2.

Figure 5. The influence of learnable σ on ZSL (**T1**) and GZSL (**H**) results (%). σ is set to 20. During the model training, the performance increases with a slight fluctuation of σ .

loss, and the unseen categories were used as a test set to verify model performance. We repeated the above operation for four times and took the average value of **T1** to ensure robustness. Table 3 shows the influence of AT with gaze loss, and we can see that AT improves the ZSL accuracy.

Effect of scaling factor σ . Figure 4 shows the results of **H** and **T1** when varying σ over $\{10, 15, 20, 25, 30\}$ under GZSL/ZSL for our model. Our method achieves the best performance on CUB and AWA2 when σ is 20. The best results are reached for SUN when σ is 25. We further explore the performance of the model when σ is a learnable parameter. Figure 5 shows the relationship between model performance and σ during the model training. We initialize σ to 20, and it changes slightly with the iteration of training. Therefore, we fixed σ to 20 for CUB and AWA2 and 25 for SUN in our experiments.

Training method analysis. An episode-based training method is used in our experiments to make the model gain better generalization ability. For each mini-batch, we sample M categories and N images for each category. We

Table 4. Influence of training method on GZSL results (%). $\mathcal{R}$ represents random sampling training method with mini-batch of 64, $\mathcal{E}$ represents episode-based training method.

Training Method	M -way	N -shot	CUB	SUN	AWA2
$\mathcal{R}$	-	-	62.3	34.9	66.4
	8	2	60.6	27.0	46.9
	8	3	64.2	27.8	61.8
	8	4	58.5	27.2	62.3
$\mathcal{E}$	12	2	62.5	33.1	68.2
	12	3	68.7	35.8	65.1
	12	4	65.5	34.5	67.4
	16	2	70.4	36.9	70.6
	16	3	68.3	36.7	68.3
	16	4	69.1	34.8	69.3

vary the value of M with $\{8, 12, 16\}$ and the value of N with $\{2, 3, 4\}$, and observe the **H** under these values. To further analyze the performance of the episode-based training method, we compare its performance with the random sampling training method with a mini-batch of 64. Table 4 shows that the episode-based training method has better performance than the random sampling training method. The model can be generalized to the recognition of all categories (seen and unseen categories) only by learning the seen categories. When $M = 16$ and $N = 2$, the model can get the highest accuracy.

4.3. GEM Evaluation

Quantitative results. We use two standard evaluation metrics for gaze estimation: Area Under the Curve (AUC) [5] and Normalized Scanpath Saliency (NSS) [43]. AUC is the area under a curve of true positive rate versus false positive rate for different thresholds on the estimated gaze map. It is a commonly used evaluation metric in saliency prediction. NSS is a simple correspondence mea-

Figure 6. Visualization of predicted gaze maps of randomly selected unseen images on CUB-VWSW from different gaze estimation methods.

Table 5. Results of different methods for gaze estimation on unseen images of CUB-VWSW dataset.

Methods	AUC	NSS
GBVS [18]	0.793	1.003
GP [21]	0.836	1.430
GEM (Ours)	0.914	2.244

sure between saliency maps and ground truth, computed as the average normalized saliency at fixated locations. We compared our GEM with two gaze estimation baselines include: bottom-up saliency prediction method Graph Based Visual Saliency (GBVS) [18] and top-down task-specific Gaze Prediction (GP) [21]. Table 5 shows the results of different methods for gaze estimation on the unseen images of CUB-VWSW dataset. We can see that our GEM achieves better gaze estimation results.

Qualitative analysis. Figure 6 shows the qualitative results of different gaze estimation methods on randomly selected unseen images from CUB-VWSW dataset. Our GEM can detect more accurate attribute regions that humans pay more attention to compared with both the bottom-up method GBVS and top-down method GP. These regions are helpful to distinguish different categories. Because our model is zero-shot learning targeted, GEM can learn more

refined category attribute areas. In the visualization results, the predicted gaze maps are more concentrated than ground truth maps. In general, these results demonstrate that the GEM plays an important role in improving the accuracy of attribute parts prediction.

5. Conclusion

In this paper, a novel goal-oriented gaze estimation method has been introduced to improve ZSL tasks. To mimic human cognitive process with unseen classes, the GEM learns the discriminative attributes with semantic query-guided attention and real human gaze (if available) supervision. Finally, the localized discriminative attributes improve the global image feature representation for ZSL. Our GEM-ZSL achieves superior or competitive performance on three ZSL benchmarks which demonstrates the effectiveness of the discriminative attributes learned by the GEM. Our work shows the promising benefits of collecting human gaze dataset and automatic gaze estimation algorithms on computer vision tasks. For the future work, further investigation on the discriminative attribute localization for ZSL is an intrinsic and important direction.

References

- [1] Radhakrishna Achanta, Sheila Hemami, Francisco Estrada, and Sabine Susstrunk. Frequency-tuned salient region detec-

- tion. In *CVPR*, pages 1597–1604, 2009. 3
- [2] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for attribute-based classification. In *CVPR*, pages 819–826, 2013. 1, 2
- [3] Zeynep Akata, Scott Reed, Daniel Walter, Honglak Lee, and Bernt Schiele. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, pages 2927–2936, 2015. 1, 2
- [4] Yashas Annadani and Soma Biswas. Preserving semantic relations for zero-shot learning. In *CVPR*, pages 7603–7612, 2018. 5, 6
- [5] Ali Borji, Hamed R Tavakoli, Dicky N Sihite, and Laurent Itti. Analysis of scores, datasets, and models in visual saliency prediction. In *ICCV*, pages 921–928, 2013. 7
- [6] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pages 177–186. Springer, 2010. 5
- [7] 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, 6
- [8] Wei-Lun Chao, Soravit Changpinyo, Boqing Gong, and Fei Sha. An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ECCV*, pages 52–68. Springer, 2016. 3, 5
- [9] 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. 5, 6
- [10] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, Philip HS Torr, and Shi-Min Hu. Global contrast based salient region detection. *PAMI*, pages 569–582, 2014. 3
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 5
- [12] Rafael Felix, Vijay BG Kumar, Ian Reid, and Gustavo Carneiro. Multi-modal cycle-consistent generalized zero-shot learning. In *ECCV*, pages 21–37, 2018. 5, 6
- [13] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *NeurIPS*, pages 2121–2129, 2013. 1, 2
- [14] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, pages 457–468, 2016. 2, 4
- [15] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *CVPR*, pages 4367–4375, 2018. 3
- [16] 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. 1
- [17] Zongyan Han, Zhenyong Fu, and Jian Yang. Learning the redundancy-free features for generalized zero-shot object recognition. In *CVPR*, pages 12865–12874, 2020. 5, 6
- [18] Jonathan Harel, Christof Koch, and Pietro Perona. Graph-based visual saliency. In *NeurIPS*, pages 545–552, 2007. 3, 8
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 5
- [20] Yifei Huang, Minjie Cai, Zhenqiang Li, Feng Lu, and Yoichi Sato. Mutual context network for jointly estimating egocentric gaze and action. *TIP*, 29:7795–7806, 2020. 3
- [21] Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Predicting gaze in egocentric video by learning task-dependent attention transition. In *ECCV*, pages 754–769, 2018. 3, 8
- [22] Dat Huynh and Ehsan Elhamifar. Fine-grained generalized zero-shot learning via dense attribute-based attention. In *CVPR*, pages 4483–4493, 2020. 5, 6
- [23] Zhong Ji, Yanwei Fu, Jichang Guo, Yanwei Pang, Zhongfei Mark Zhang, et al. Stacked semantics-guided attention model for fine-grained zero-shot learning. In *NeurIPS*, pages 5995–6004, 2018. 2
- [24] Huajie Jiang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Transferable contrastive network for generalized zero-shot learning. In *ICCV*, pages 9765–9774, 2019. 5, 6
- [25] Ming Jiang, Shengsheng Huang, Juanyong Duan, and Qi Zhao. Salicon: Saliency in context. In *CVPR*, pages 1072–1080, 2015. 3
- [26] Nour Kaessli, Zeynep Akata, Bernt Schiele, and Andreas Bulling. Gaze embeddings for zero-shot image classification. In *CVPR*, pages 4525–4534, 2017. 1, 6
- [27] Rohit Keshari, Richa Singh, and Mayank Vatsa. Generalized zero-shot learning via over-complete distribution. In *CVPR*, pages 13300–13308, 2020. 5, 6
- [28] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, pages 10215–10224, 2018. 1
- [29] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1
- [30] Vinay Kumar Verma, Gundeep Arora, Ashish Mishra, and Piyush Rai. Generalized zero-shot learning via synthesized examples. In *CVPR*, pages 4281–4289, 2018. 1
- [31] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, pages 951–958, 2009. 1
- [32] Kai Li, Martin Renqiang Min, and Yun Fu. Rethinking zero-shot learning: A conditional visual classification perspective. In *ICCV*, pages 3583–3592, 2019. 3, 6
- [33] Yan Li, Junge Zhang, Jianguo Zhang, and Kaiqi Huang. Discriminative learning of latent features for zero-shot recognition. In *CVPR*, pages 7463–7471, 2018. 5, 6
- [34] Shichen Liu, Mingsheng Long, Jianmin Wang, and Michael I Jordan. Generalized zero-shot learning with deep calibration network. In *NeurIPS*, pages 2005–2015, 2018. 2
- [35] Yang Liu, Jishun Guo, Deng Cai, and Xiaofei He. Attribute attention for semantic disambiguation in zero-shot learning. In *ICCV*, pages 6698–6707, 2019. 2, 5, 6
- [36] Yang Liu, Lei Zhou, Xiao Bai, Lin Gu, Tatsuya Harada, and Jun Zhou. Information bottleneck constrained latent bidi-

- rectional embedding for zero-shot learning. *arXiv preprint arXiv:2009.07451*, 2020. 1
- [37] Chunjie Luo, Jianfeng Zhan, Xiaohe Xue, Lei Wang, Rui Ren, and Qiang Yang. Cosine normalization: Using cosine similarity instead of dot product in neural networks. In *ICANN*, pages 382–391, 2018. 3
- [38] Shaobo Min, Hantao Yao, Hongtao Xie, Chaoqun Wang, Zheng-Jun Zha, and Yongdong Zhang. Domain-aware visual bias eliminating for generalized zero-shot learning. In *CVPR*, pages 12664–12673, 2020. 5, 6
- [39] Pedro Morgado and Nuno Vasconcelos. Semantically consistent regularization for zero-shot recognition. In *CVPR*, pages 6060–6069, 2017. 1, 2
- [40] Junting Pan, Elisa Sayrol, Xavier Giro-i Nieto, Kevin McGuinness, and Noel E O’Connor. Shallow and deep convolutional networks for saliency prediction. In *CVPR*, pages 598–606, 2016. 3
- [41] Genevieve Patterson and James Hays. Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*, pages 2751–2758, 2012. 5
- [42] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014. 4
- [43] Robert J Peters, Asha Iyer, Laurent Itti, and Christof Koch. Components of bottom-up gaze allocation in natural images. *Vision research*, pages 2397–2416, 2005. 7
- [44] Bernardino Romera-Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, pages 2152–2161, 2015. 1, 2
- [45] 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. 1, 5, 6
- [46] Yuming Shen, Jie Qin, Lei Huang, Li Liu, Fan Zhu, and Ling Shao. Invertible zero-shot recognition flows. In *ECCV*, pages 614–631, 2020. 1, 5, 6
- [47] Jie Song, Chengchao Shen, Yezhou Yang, Yang Liu, and Mingli Song. Transductive unbiased embedding for zero-shot learning. In *CVPR*, pages 1024–1033, 2018. 2, 5, 6
- [48] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018. 5, 6
- [49] Tristan Sylvain, Linda Petrini, and Devon Hjelm. Locality and compositionality in zero-shot learning. In *ICLR*, 2020. 1
- [50] Anne M Treisman and Garry Gelade. A feature-integration theory of attention. *Cognitive psychology*, 12(1):97–136, 1980. 3
- [51] Maunil R Vyas, Hemanth Venkateswara, and Sethuraman Panchanathan. Leveraging seen and unseen semantic relationships for generative zero-shot learning. In *ECCV*, pages 70–86, 2020. 5, 6
- [52] Wenguan Wang, Jianbing Shen, Jianwen Xie, Ming-Ming Cheng, Haibin Ling, and Ali Borji. Revisiting video saliency prediction in the deep learning era. *PAMI*, 2019. 3
- [53] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010. 5
- [54] 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. 2
- [55] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *PAMI*, pages 2251–2265, 2018. 5
- [56] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *CVPR*, pages 5542–5551, 2018. 1, 5, 6
- [57] Yongqin Xian, Bernt Schiele, and Zeynep Akata. Zero-shot learning—the good, the bad and the ugly. In *CVPR*, pages 4582–4591, 2017. 2
- [58] 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. 1, 2, 5, 6
- [59] Guo-Sen Xie, Li Liu, Fan Zhu, Fang Zhao, Zheng Zhang, Yazhou Yao, Jie Qin, and Ling Shao. Region graph embedding network for zero-shot learning. In *ECCV*, pages 562–580, 2020. 1, 2, 3, 5, 6
- [60] Wenjia Xu, Yongqin Xian, Jiuniu Wang, Bernt Schiele, and Zeynep Akata. Attribute prototype network for zero-shot learning. In *NeurIPS*, 2020. 1, 2, 3, 5, 6
- [61] Hyeonwoo Yu and Beomhee Lee. Zero-shot learning via simultaneous generating and learning. In *NeurIPS*, pages 46–56, 2019. 1
- [62] Yunlong Yu, Zhong Ji, Jungong Han, and Zhongfei Zhang. Episode-based prototype generating network for zero-shot learning. In *CVPR*, pages 14035–14044, 2020. 5, 6
- [63] Zhou Yu, Jun Yu, Jianping Fan, and Dacheng Tao. Multi-modal factorized bilinear pooling with co-attention learning for visual question answering. In *ICCV*, pages 1821–1830, 2017. 2, 4
- [64] Li Zhang, Tao Xiang, and Shaogang Gong. Learning a deep embedding model for zero-shot learning. In *CVPR*, pages 2021–2030, 2017. 2
- [65] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for fine-grained image recognition. In *ICCV*, pages 5209–5217, 2017. 4
- [66] Pengkai Zhu, Hanxiao Wang, and Venkatesh Saligrama. Generalized zero-shot recognition based on visually semantic embedding. In *CVPR*, pages 2995–3003, 2019. 1
- [67] Yizhe Zhu, Jianwen Xie, Zhiqiang Tang, Xi Peng, and Ahmed Elgammal. Semantic-guided multi-attention localization for zero-shot learning. In *NeurIPS*, pages 14943–14953, 2019. 1, 2, 5, 6