On Implicit Attribute Localization for Generalized Zero-Shot Learning

Shiqi Yang, Kai Wang, Luis Herranz and Joost van de Weijer

Abstract—Zero-shot learning (ZSL) aims to discriminate images from unseen classes by exploiting relations to seen classes via their attribute-based descriptions. Since attributes are often related to specific parts of objects, many recent works focus on discovering discriminative regions. However, these methods usually require additional complex part detection modules or attention mechanisms. In this paper, 1) we show that common ZSL backbones (without explicit attention nor part detection) can implicitly localize attributes, yet this property is not exploited. 2) Exploiting it, we then propose SELAR, a simple method that further encourages attribute localization, surprisingly achieving very competitive generalized ZSL (GZSL) performance when compared with more complex state-of-the-art methods. Our findings provide useful insight for designing future GZSL methods, and SELAR provides an easy to implement yet strong baseline.

Index Terms—Zero-shot learning, Attribute Localization

I. INTRODUCTION

Visual classification with deep convolutional neural networks has achieved remarkable success [1], [2], even surpassing humans on some benchmarks [3]. This success, however, requires enough training images per class (tens to hundreds of images), which is often not the case in practice. The visual data to learn new classes may be scarce (i.e. few-shot learning) or inexistent (i.e. zero-shot learning -ZSL-). Humans, in contrast, are able to infer new classes from few or even no visual examples, just from a semantic description that connects them to known concepts . Thus, ZSL is a desirable capability for computer vision systems, allowing them to recognize a larger set of classes via their semantic descriptions.

The most common representations in zero shot learning are global visual features extracted from a pretrained feature extractor, which are readily available off-the-shelf from previous works [4], [5]. These global visual features are then projected to a semantic space [6], [7] or to an intermediate space [8], where the comparison with semantic representations takes place. In this paper, we focus on generalized zero-shot learning (GZSL), where the test set includes both seen and unseen classes. A major problem in GZSL is the model bias towards seen classes. Existing works can be roughly divided into two lines: generative and non-generative. Non-generative methods [6], [7], [8], [4], [5], [9], [10], [11], [12], [13], [14], [15], [16] focus on designing and learning a good visual-semantic alignment. Generative methods [17], [18], [19], [20],

[21], [22], [23] adopt generative models to synthesize visual features of unseen classes. Generative methods have much higher performance than non-generative methods, since the classifier is trained with synthetic visual features of both seen and unseen classes (real and synthetic features, respectively), thus successfully avoiding the prediction bias towards seen classes. Note that generative methods usually need access to the descriptions of unseen classes during training (however a counter-example is [22]), while non-generative GZSL methods generally assume they do not have access. Thus, a comparison between these two settings is not fair.

Most non-generative approaches focus on the role of the classifier and the semantic models, directly relying on global representations extracted by a pretrained network. The potential of local representations has been explored mainly recently in two directions: part detection [9], [10], and attention mechanisms [11], [12], [13]. [16] takes a step further, and combines part attention and leverages graph convolutional networks to reason about the parts relations. However, training part detectors normally requires additional and expensive annotation data (i.e. part ids, bounding boxes) to train the part detector. In contrast, attention mechanisms focus on discovering discriminative regions without requiring explicit region annotations. Both part detectors and attention mechanisms are significantly more complex and arguably more difficult to train than our proposed approach. Their representations are also essentially different to ours, since the attributes are not localized separately, but only a few regions are extracted. Besides of those, [15] tries to find a more discriminative latent attribute space, and [14] learns another label space where labels of unseen classes can be regarded as linear combination of labels of seen classes, which is more suitable for the GZSL. However, none of these methods notices the potential for attribute localization implicit in the feature extractor itself.

In this paper, we show a simple pipeline with fundamental modules is actually localizing attributes implicitly, a fact that is overlooked in literature. Then we show that the specific spatial aggregation method (implicitly, global average pooling - GAP - in most existing methods) becomes a critical choice in the design, since it can influence the degree of localization of the attributes. We then propose SELAR, a simple yet powerful variant of the previous baseline that uses global maximum pooling (GMP) as aggregation method. This simple modification further encourages attribute localization, boosting the performance and outperforming most GZSL methods with explicit attention. This suggests that implicit attribute localization provides better and more efficient representations than those learned with complex explicit attention modules.
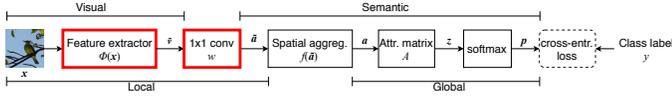
Fig. 1: A simple pipeline. Trainable modules are highlighted in red. The final fully connected layer is initialized with attribute matrix A and fixed.
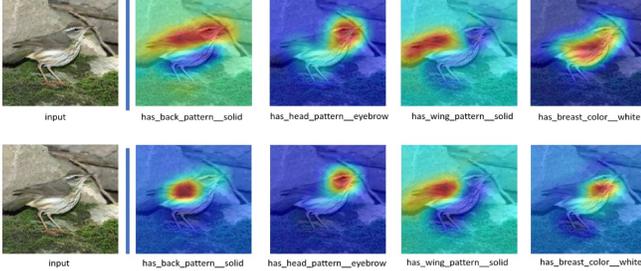
Fig. 2: Attribute Activation Maps (AAMs) output by 1x1 conv from baseline (top) and SELAR (bottom) on CUB.

II. METHOD

A. Task Definition

In the GZSL, the training set contains seen classes and is defined as $\mathcal{S} \equiv \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^{N_s}$, where i denotes the i-th image of the seen class and $y_i^s \in \mathcal{Y}^\mathcal{S}$ is its class label. The test set contains unseen classes and is defined as $\mathcal{U} \equiv \{(\mathbf{x}_j^u, y_j^u)\}_{j=1}^{N_u}$. The sets of seen and unseen classes are disjoint, i.e. $\mathcal{Y}^\mathcal{S} \cap \mathcal{Y}^\mathcal{U} = \emptyset$. The semantic information about a particular class y is obtained by the class embedding function as $\psi(y_i)$. In the case of attribute-based representations with L attributes, the class prototype $\psi_i = \psi(y_i)$ is simply a L-dimensional (binary or real valued) attribute vector . In this way, the semantic information about all seen classes can be conveniently captured in a $|\mathcal{Y}^\mathcal{S}| \times L$-dimensional attribute matrix $A^\mathcal{S} \equiv [\psi_1, \ldots, \psi_{|\mathcal{Y}^\mathcal{S}|}]^\top$. Similarly, for unseen classes we obtain $A^\mathcal{U} \equiv [\psi_1, \ldots, \psi_{|\mathcal{Y}^\mathcal{U}|}]^\top$. Finally, the evaluation in the GZSL setting considers a test set that includes both seen and unseen classes, i.e. $\mathcal{Y}^{\mathcal{S}\mathcal{U}} = \mathcal{Y}^\mathcal{S} \cup \mathcal{Y}^\mathcal{U}$.

B. Implicit Attribute Localization in ZSL

We formulate ZSL as a classification problem, using a deep convolutional neural network (CNN) that internally projects visual features to the semantic space and is trained end-to-end with cross-entropy loss on seen data. In particular, we are interested in certain intermediate representations: the *local visual feature* $\tilde{\mathbf{v}} \in \mathbb{R}^{M \times M \times D}$, the *global visual feature* $\mathbf{v} \in \mathbb{R}^D$ which is output by spatial aggregation, the *global semantic feature* $\mathbf{a} \in \mathbb{R}^L$, and the logits or *unnormalized class-scores* $\mathbf{z} \in \mathbb{R}^{|\mathcal{Y}^\mathcal{S}|}$. These intermediate representations lie in three distinctive spaces: the D-dimensional visual space, the L-dimensional semantic space (where L is the number of attributes in our case) and the $|\mathcal{Y}^\mathcal{S}|$-dimensional class space. For convenience, we can split the deep network into several modules: the *feature extractor* $\tilde{\mathbf{v}} = \phi(\mathbf{x})$, the *spatial aggregation* operation $\mathbf{v} = f(\tilde{\mathbf{v}})$, and the linear projection to the semantic space $\mathbf{a} = W\mathbf{v}$, parametrized by the embedding (or mapping) matrix $W \in \mathbb{R}^{L \times D}$, i.e. a fully connected layer. The embedding matrix is trainable from scratch, while the

feature extractor is usually pretrained and can be optionally fine-tuned. Finally, the overall loss to minimize is

$$\mathcal{L} = \mathbb{E}_{(\mathbf{x},y) \sim \mathcal{S}} \left[\mathcal{CE} \left(\text{softmax} \left(A^\mathcal{S} W f(\phi(\mathbf{x})) \right), y \right) \right] \quad (1)$$

where $\mathcal{CE}$ is the cross-entropy loss. The $A^\mathcal{S} W f(\phi(\mathbf{x}))$ is known as bilinear compatibility score. It measures the similarity between the attribute vector and embedded visual features.

Normally in ZSL, the feature extractor is a pretrained VGG [2] or ResNet [1] network, where spatial aggregation is global average pooling (GAP). Since both are linear operations, we can switch the order of GAP and the linear layer W (which now becomes a 1×1 convolution), as shown in Fig. 1. We emphasize that this simple pipeline (denoted as *baseline*) already includes the key elements of ZSL methods operating on local features, that is, feature extractor (with the spatial aggregation) and embedding/mapping layer, since most methods at least contain these modules.

Since the baseline in Fig. 1 (with GAP) is identical to a normal classification pipeline, we can get the Class Activation Map (CAM) [24]. The CAM is computed as:

$$C_c = \sum_{i=1}^{N_{att}} A_{c,i} \cdot \tilde{a}_i \quad (2)$$

where the C_c is the CAM of class c. It shows the most contributing region for the current class prediction. N_{att} is the number of attributes, and $A_{c,i}$ means the i-th value in the attribute vector A_c. And the $\tilde{a}_i$ is the i-th channel of the feature map which is output by the feature extractor in Fig. 1.

CAM shows which part of the image contributes more to the current class prediction. The attribute vector A_c gives the importance of all attributes for the class c. It actually weighs the different feature maps $\tilde{a}_i$ to achieve object localization. It is reasonable to conjecture that each feature map $\tilde{a}_i$ is related to a specific attribute (the $i-$th attribute in the attribute vector). When we visualize the four feature maps (referred to as Attribute Activation Map, AAM) output by the feature extractor from the baseline, as shown in the upper part of Fig. 2, those four AAMs correspond to the four attribute which have the highest values in the attribute vector, *i.e.*, having a higher contribution to the current class. We can observe that the AAMs are quite accurate performing weakly supervised attribute localization. Thus, the visual-semantic embedding is actually implicitly localizing the attribute regions.

C. SELAR: a simple yet effective method

In Figs. 3a and 3b, we first visualize the embedded visual vectors (for both seen and unseen classes) from the baseline on the CUB dataset (312 dimension corresponding to 312 attributes) for 300 images (50 images per class). Each row is one embedded visual vector. We put together those vectors from the same class and join them with the attribute vector of this specific class (depicted in red). Because those features are already embedded into the semantic space, they are expected to be similar to the attribute vector. Ideally, the discriminative features are expected to have high values only for those attributes with high values in the ground-truth class attribute vector. But observing Figs. 3a and 3b, embedded visual feature

| (a) Baseline (seen) | (b) Baseline (unseen) | (c) SELAR (seen) | (d) SELAR (unseen) |
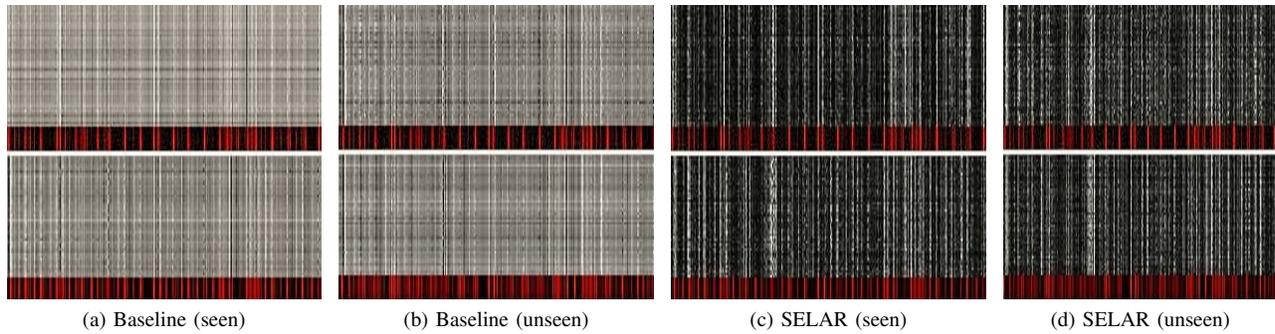
Fig. 3: Global semantic representations (rows) of 50 images per class (randomly selected) of 2 classes (super-rows). Each column corresponds to one of the 312 attributes. The description corresponding to the class (attribute vector) is shown in red (scaled to multiple rows). **The ideal features (white one) should only have similar and sparse activation as the attribute vectors (red one).**

vectors are also activated for some non-existing attributes (the ones which are not red in the attribute vector).

We posit that this is due to the fact that GAP aggregates all spatial information of the feature maps (i.e. AAMs). Each spatial location in AAMs already corresponds to a large receptive field, while most attributes are typically localized in a particular region, thus producing noisy features from GAP.

In order to make the embedded visual vector more sparse and discriminative, we focus on obtaining more accurate AAMs. Therefore, we replace the GAP with global max pooling (GMP). The rationale behind this choice is that GMP has to focus on a single spatial location of the input AAMs (instead of averaging them as in GAP), thus enforcing the model to learn AAMs that are more localized . The framework is shown in Fig. 1. We refer to this approach as *SELAR* (**S**imple and **E**ffective **L**ocalized **A**ttribute **R**epresentation). Compared with the baseline, SELAR produces AAMs which correspond to more relevant regions, as shown in Fig. 2. And the visualization of embedded visual vector shown in Figs. 3c and 3d clearly shows that SELAR generates more sparse and discriminative feature vectors, resembling more the actual attribute vectors describing the class.

III. EXPERIMENTS

Settings. We evaluate our method on three datasets: CUB [30], SUN [31] and AWA2 [4] under the challenging GZSL setting. We denote the accuracy on unseen classes and seen classes as Acc_U and Acc_S, respectively, and the evaluation metric for GZSL is the harmonic mean, calculated as $H = 2 * Acc_U * Acc_S/(Acc_U + Acc_S)$. Our model is built on top of ImageNet pretrained model. We initialize the last fully connected layer with the L2-normalized attribute matrix, then remaining fixed. For fair comparison, we distinguish between methods that assume the descriptions of unseen classes are unknown during training (typically non-generative methods), and methods that assume they are known (generative methods).

GZSL performance. Table I shows the results. Regarding methods that do not assume knowledge about unseen classes, SELAR achieves state-of-the-art performance on CUB and competitive results on SUN and AWA2. Notably, it surpasses other more complex methods on the fine-grained CUB dataset

significantly. We conjecture that this is likely due to the fact that attributes of CUB are very localized, and easily linked to visual patterns, unlike in AWA2 and SUN, which have abstract attributes such as *smart* and *domestic*. Among the methods in Table I, AREN [11], JLA [28] and AttentionZSL [29] utilize extra modules to localize parts or obtain attention maps. AREN also has one branch with explicit self-attention and GMP, and SGMA [12] has additional part detection modules. However, these methods obtain inferior results on most datasets compared to our method, despite of their increased complexity. While SGMA achieves state-of-the-art performance on AWA2, however, it requires four forward passes through the feature extractor and larger input resolution. Specifically for SUN, using GAP or GMP does not make much difference. We posit that this is due to the fact that attributes in the SUN dataset are not always clearly localizable (like the attribute *natural light*); whether to consider all these regions by GAP or only a single location by GMP does not influence a lot. We also report the average H over all these datasets (we do not include here methods with results only on two datasets) in Table I. For this metric SELAR has the highest value among all these methods. The ratio S/U (only for non-generative and without calibration) suggests that our method has lower seen-unseen bias than others. In general, we observe that two simple modifications of the backbone network (i.e. mapping to localized semantic space with a 1×1 convolution and encouraging more discriminative localization with GMP) result in significant gains and lower bias, and surprisingly good results, especially considering that the other methods use much more complex architectures.

We also report results for SELAR with calibration stacking [32]. Calibration stacking helps to avoid bias towards seen classes by decreasing the prediction score for all seen classes during evaluation. However, finding the precise factor requires observing a small set of validation images (both seen and useen). In addition, for reference, we also include results of several generative models that assume knowledge (except IZF-NBC [22]) of the descriptions of unseen classes (i.e. images are unseen, but class descriptions are seen), which synthesize features for unseen data, and thus training a less biased classifier with seen and unseen classes.

Method	SUN				CUB				AWA2				$\bar{H}$
	U	S	H	S/U	U	S	H	S/U	U	S	H	S/U	
Non-generative methods (Unseen class descriptions are unknown during training)													
PSR [25] *w/o*	20.8	37.2	26.7	1.79	24.6	54.3	33.9	2.21	20.7	73.8	32.3	3.57	31.0
DCN [26] *w/o*	**25.5**	37.0	**30.2**	1.45	28.4	60.7	38.7	2.14	25.5	84.2	39.1	3.30	36.0
MIIR [27] *w/o*	22.0	34.1	26.7	1.55	30.4	65.8	41.2	2.16	17.6	87.0	28.9	4.94	32.3
AREN [11]	19.0	38.8	25.5	2.04	38.9	78.7	52.1	2.02	17.5	93.2	29.5	5.33	35.7
JLA [28]	23.2	36.6	28.4	1.58	36.6	59.8	45.4	1.63	24.5	91.6	38.3	3.74	37.4
AttentionZSL [29]	18.5	**40.0**	25.3	2.16	36.2*	**80.9***	50.0*	2.23	27.0*	**93.4***	41.9*	3.46	39.1
SGMA [12]	–	–	–	–	36.7*	71.3*	48.5*	1.94	**37.6***	87.1*	**52.5***	**2.32**	–
Baseline *w/o*	23.8	32.0	27.3	**1.34**	32.1	63.0	42.5	1.96	12.0	87.2	21.0	7.27	30.3
Baseline	23.4	37.2	28.7	1.59	39.0(37.1*)	74.2(73.2*)	51.1(49.2*)	1.97	13.7(14.6*)	90.4(77.0*)	23.8(24.5*)	6.60	33.9
SELAR *w/o*	22.8	31.6	26.5	1.39	43.5	54.0	48.2	1.64	31.6	80.3	45.3	2.54	41.9
SELAR	23.8	37.2	29.0	1.56	43.0(**51.4***)	76.3(75.2*)	55.0(**61.0***)	1.46	32.9(29.5*)	78.7(80.2*)	46.4(43.2*)	2.39	**45.5**
Generative methods (Unseen class descriptions are known during training)													
f-CLSWGAN [17]	42.6	36.6	39.4	–	43.7	57.7	49.7	–	52.1	68.9	59.4	–	49.5
f-VAEGAN-D2 [18]	45.1	38.0	41.3	–	48.4	60.1	53.6	–	57.6	70.6	63.5	–	52.8
CADA-VAE [19]	**47.2**	35.7	40.6	–	51.6	53.5	52.4	–	55.8	75.0	63.9	–	52.3
attr-BImag [20]	22.4	40.1	29.2	–	41.3	**77.7**	53.9	–	**60.0**	73.8	**66.2**	–	49.8
class-attr-BImag [20]	21.7	40.7	28.2	–	44.1	73.6	55.9	–	51.4	**76.9**	61.6	–	48.6
‡IZF-NBC [22]	44.5	**50.6**	47.4		44.2	56.3	49.5		58.1	76.0	65.9		**54.3**
†**SELAR**	40.5	32.9	36.3	–	**62.4**	64.9	**63.6**	–	52.0	71.9	60.3	–	53.4

TABLE I: Generalized Zero-Shot Learning on Proposed Split (PS). **U** = Top-1 accuracy on $\mathcal{Y}_\mathcal{U}$, **S** = Top-1 accuracy on $\mathcal{Y}_\mathcal{S}$, H = harmonic mean, **S/U** can show the bias towards seen class, $\bar{H}$ denotes the average over the H on three datasets. * indicates results using VGG19 as feature extractor while others use ResNet-101. *w/o* means not finetuning feature extractor. We highlight the best results for both non-generative and generative methods. † indicates using calibration, which assumes access to a small validation set of seen and unseen classes. ‡ means not requiring accessing to unseen class descriptions.

Type	Space	U	S	H
GAP	visual, attribute, class (all equivalent)	37.1	73.2	49.2
GMP	visual	39.1	**80.7**	52.7
GMP	attribute (ours)	**51.4**	75.2	**61.0**
GMP	class	26.3	74.7	38.9

TABLE II: Ablation study of pooling operations and pooling spaces. The results are reported on the CUB dataset.
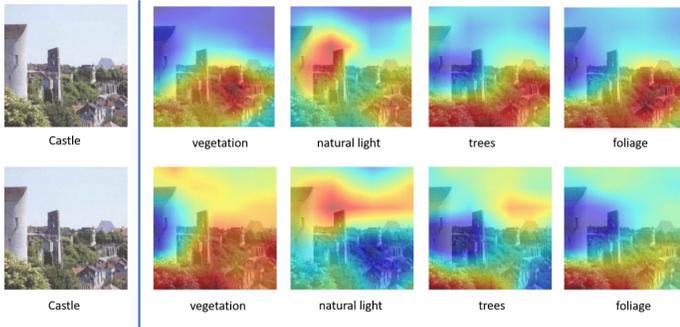
Fig. 4: Visualization of attribute maps on SUN from the baseline (top one) and SELAR (bottom one). Below the image are the corresponding attributes.

Aggregation method and aggregation space. We investigate the optimal location to aggregate local features into global ones. The ablation study shown in Table II evaluates GAP and GMP in three different spaces: visual (before the em-
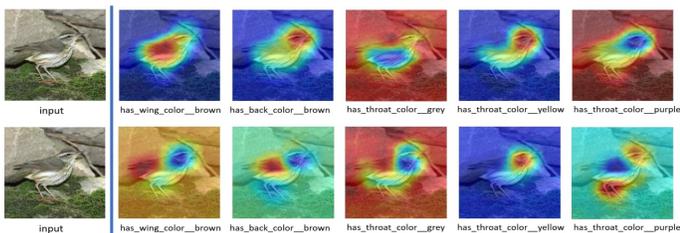
Fig. 5: Visualization of attribute maps on CUB from baseline (upper part) and SELAR (lower part). Those attributes have **lower** value in the attribute vector.

bedding layer), attribute (after embedding layer) and class (after attribute mapping), which also correspond to the order in which features are mapped to the different spaces in our classification pipeline. Note that GAP in either of the three spaces is equivalent due to linearity. Results suggest that the best performance is achieved by GMP in the attribute space.

Attribute Activation Maps. Here we visualize AAMs $\tilde{a}$, for both the baseline and SELAR on SUN in Fig. 4. Whereas on the CUB and AWA datasets the attributes are present in clearly localizable small regions, for the SUN dataset this is not the case. This maybe the reason why SELAR and the baseline have similar performance on SUN. While we cannot guarantee that each AAM corresponds to the true attribute, the network learns to correlate the related region with its activation value automatically. We find that AAMs with high activation values (>85) are always highly related to the corresponding attribute. We show additional AAMs with lower attribute values in Fig. 5, those attributes have either middle value or do not exist (0). In those cases, the AAM sometimes corresponds to a specific attribute, but sometimes not.

IV. CONCLUSIONS

In this paper, we highlight an overlooked fact in zero-shot learning research, that is, that the feature extractor backbones of common ZSL pipelines can implicitly localize attributes without any explicit region localization module such as attention or part detection. Based on this finding, we propose SELAR, a simple GZSL framework that efficiently extracts highly discriminative representations based on localized attributes. Our method is surprisingly effective, achieving competitive results compared to much more complex state-of-the-art methods. Our findings provide useful insight into how to further design effective and efficient (G)ZSL methods, and advocate SELAR as a strong, yet simple and easy to implement, baseline for future zero-shot learning research.

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[3] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.

[4] Y. Xian, C. H. Lampert, B. Schiele, and Z. Akata, "Zero-shot learning - a comprehensive evaluation of the good, the bad and the ugly," *TPAMI*, pp. 2251–2265, 2018.

[5] L. Zhang, T. Xiang, and S. Gong, "Learning a deep embedding model for zero-shot learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2021–2030.

[6] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, "Label-embedding for image classification," *IEEE TPAMI*, vol. 38, no. 7, pp. 1425–1438, 2016.

[7] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, M. A. Ranzato, and T. Mikolov, "Devise: A deep visual-semantic embedding model," in *NIPS*, 2013, pp. 2121–2129.

[8] Z. Zhang and V. Saligrama, "Zero-shot learning via semantic similarity embedding," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4166–4174.

[9] M. Elhoseiny and M. Elfeki, "Creativity inspired zero-shot learning," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 5784–5793.

[10] Y. Zhu, M. Elhoseiny, B. Liu, X. Peng, and A. Elgammal, "A generative adversarial approach for zero-shot learning from noisy texts," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1004–1013.

[11] G.-S. Xie, L. Liu, X. Jin, F. Zhu, Z. Zhang, J. Qin, Y. Yao, and L. Shao, "Attentive region embedding network for zero-shot learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[12] Y. Zhu, J. Xie, Z. Tang, X. Peng, and A. Elgammal, "Semantic-guided multi-attention localization for zero-shot learning," in *Advances in Neural Information Processing Systems*, 2019, pp. 14 917–14 927.

[13] D. Huynh and E. Elhamifar, "Fine-grained generalized zero-shot learning via dense attribute-based attention," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4483–4493.

[14] Y. Liu, X. Gao, Q. Gao, J. Han, and L. Shao, "Label-activating framework for zero-shot learning," *Neural Networks*, vol. 121, pp. 1–9, 2020.

[15] H. Jiang, R. Wang, S. Shan, Y. Yang, and X. Chen, "Learning discriminative latent attributes for zero-shot classification," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 4223–4232.

[16] G.-S. Xie, L. Liu, F. Zhu, F. Zhao, Z. Zhang, Y. Yao, J. Qin, and L. Shao, "Region graph embedding network for zero-shot learning," in *European Conference on Computer Vision*. Springer, 2020, pp. 562–580.

[17] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata, "Feature generating networks for zero-shot learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 5542–5551.

[18] Y. Xian, S. Sharma, B. Schiele, and Z. Akata, "f-vaegan-d2: A feature generating framework for any-shot learning," in *IEEE Computer Vision and Pattern Recognition (CVPR)*, 2019.

[19] E. Schonfeld, S. Ebrahimi, S. Sinha, T. Darrell, and Z. Akata, "Generalized zero-and few-shot learning via aligned variational autoencoders," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8247–8255.

[20] K. Wang, L. Herranz, A. Dutta, and J. van de Weijer, "Bookworm continual learning: beyond zero-shot learning and continual learning," *arXiv preprint arXiv:2006.15176*, 2020.

[21] X. Song, H. Zeng, S. Zhang, L. Herranz, and S. Jiang, "Generalized zero-shot learning with multi-source semantic embeddings for scene recognition," in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 3976–3985.

[22] Y. Shen, J. Qin, L. Huang, L. Liu, F. Zhu, and L. Shao, "Invertible zero-shot recognition flows," in *European Conference on Computer Vision*. Springer, 2020, pp. 614–631.

[23] R. Gao, X. Hou, J. Qin, J. Chen, L. Liu, F. Zhu, Z. Zhang, and L. Shao, "Zero-vae-gan: Generating unseen features for generalized and trans-ductive zero-shot learning," *IEEE Transactions on Image Processing*, vol. 29, pp. 3665–3680, 2020.

[24] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.

[25] Y. Annadani and S. Biswas, "Preserving semantic relations for zero-shot learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7603–7612.

[26] S. Liu, M. Long, J. Wang, and M. I. Jordan, "Generalized zero-shot learning with deep calibration network," in *Advances in Neural Information Processing Systems*, 2018, pp. 2005–2015.

[27] Y. L. Cacheux, H. L. Borgne, and M. Crucianu, "Modeling inter and intra-class relations in the triplet loss for zero-shot learning," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 10 333–10 342.

[28] Y. Li and D. Wang, "Joint learning of attended zero-shot features and visual-semantic mapping," in *BMVC*, 2019.

[29] Y. Liu, J. Guo, D. Cai, and X. He, "Attribute attention for semantic disambiguation in zero-shot learning," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2019.

[30] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona, "Caltech-UCSD Birds 200," California Institute of Technology, Tech. Rep. CNS-TR-2010-001, 2010.

[31] G. Patterson, C. Xu, H. Su, and J. Hays, "The sun attribute database: Beyond categories for deeper scene understanding," *International Journal of Computer Vision*, vol. 108, no. 1-2, pp. 59–81, 2014.

[32] W.-L. Chao, S. Changpinyo, B. Gong, and F. Sha, "An empirical study and analysis of generalized zero-shot learning for object recognition in the wild," in *European Conference on Computer Vision*. Springer, 2016, pp. 52–68.