

Multi-Head Self-Attention via Vision Transformer for Zero-Shot Learning

Faisal Alamri and Anjan Dutta

Department of Computer Science, University of Exeter, United Kingdom

Abstract

Zero-Shot Learning (ZSL) aims to recognise unseen object classes, which are not observed during the training phase. The existing body of works on ZSL mostly relies on pretrained visual features and lacks the explicit attribute localisation mechanism on images. In this work, we propose an attention-based model in the problem settings of ZSL to learn attributes useful for unseen class recognition. Our method uses an attention mechanism adapted from Vision Transformer to capture and learn discriminative attributes by splitting images into small patches. We conduct experiments on three popular ZSL benchmarks (i.e., AWA2, CUB and SUN) and set new state-of-the-art harmonic mean results on all the three datasets, which illustrate the effectiveness of our proposed method.

Keywords: Generalised zero-shot learning, Inductive learning, Attention, Semantic embedding, Vision Transformer.

1 Introduction

Relying on massive annotated datasets, significant progress has been made on many visual recognition tasks, which is mainly due to the widespread use of different deep learning architectures [Ren et al., 2015, Dosovitskiy et al., 2021, Khan et al., 2021]. Despite these advancements, recognising any arbitrary real-world object still remains a daunting challenge as it is unrealistic to label all the existing object classes on the earth. Zero-Shot Learning (ZSL) addresses this problem, requiring images from the *seen* classes during the training, but has the capability of recognising *unseen* classes during the inference [Xian et al., 2019a, Xie et al., 2019, Xu et al., 2020, Federici et al., 2020]. Here the central insight is that all the existing categories share a common semantic space and the task of ZSL is to learn a mapping from the imagery space to the semantic space with the help of side information (attributes, word embeddings) [Xian et al., 2017, Mikolov et al., 2013, Pennington et al., 2014] available with the seen classes during the training phase so that it can be used to predict the class information for the unseen classes during the inference time.

Most of the existing ZSL methods [Xian et al., 2018, Schönfeld et al., 2019] depends on pretrained visual features and necessarily focus on learning a compatibility function between the visual features and semantic attributes. Although modern neural network models encode local visual information and object parts [Xie et al., 2019], they are not sufficient to solve the localisation issue in ZSL models. Some attempts have also




Figure 1: Our method embeds each attribute-based feature with the semantic space. It learns the visual discriminative features through multi-head attention. *Best to view in colour*: colours in the image correspond to the same-colour attribute in the semantic space.

been made by learning visual attention that focuses on some object parts [Zhu et al., 2019]. However, designing a model that can exploit a stronger attention mechanism is relatively unexplored.

Therefore, to alleviate the above shortcomings of visual representations in ZSL models, in this paper, we propose a Vision Transformer (ViT) [Dosovitskiy et al., 2021] based multi-head self-attention model for solving the ZSL task. Our main contribution is to introduce ViT for enhancing the visual feature localisation to solve the zero-shot learning task. Without any object part-level annotation or detection, this is the first attempt to introduce ViT into ZSL. As illustrated in Figure 1, our method maps the visual features of images to the semantic space with the help of scaled dot-product of multi-head attention employed in ViT. We have also performed detailed experimentation on three public datasets (i.e., AWA2, CUB and SUN) following Generalised Zero-Shot Learning (GZSL) setting and achieved very encouraging results on all of them, including the new state-of-the-art harmonic mean on all the datasets.

2 Related Work

Zero-Shot Learning: ZSL is employed to bridge the gap between seen and unseen classes using semantic information, which is done by computing similarity function between visual features and previously learned knowledge [Romera-Paredes and Torr, 2015]. Various approaches address the ZSL problem by learning probabilistic attribute classifiers to predict class labels [Lampert et al., 2009, Norouzi et al., 2014] and by learning linear [Frome et al., 2013, Akata et al., 2015, Akata et al., 2016], and non-linear [Xian et al., 2016] compatibility function associating image features and semantic information. Recently proposed generative models synthesise visual features for the unseen classes [Xian et al., 2018, Schönfeld et al., 2019]. Although those models achieve better performances compared to classical models, they rely on features of trained CNNs. Recently, attention mechanism is adapted in ZSL to integrate discriminative local and global visual features. Among them, S²GA [Yu et al., 2018] and AREN [Xie et al., 2019] use an attention-based network with two branches to guide the visual features to generate discriminative regions of objects. SGMA [Zhu et al., 2019] also applies attention to jointly learn global and local features from the whole image and multiple discovered object parts. Very recently, APN [Xu et al., 2020] proposes to divide an object into eight groups and learns a set of attribute prototypes, which further help the model to decorrelate the visual features. Partly inspired by the success of attention-based models, in this paper, we propose to learn local and global features using multi-scaled-dot-product self-attention via the Vision Transformer model, which to the best of our knowledge, is the first work on ZSL involving Vision Transformer. In this model, we employ multi-head attention after splitting the image into fixed-size patches so that it can attend to each patch to capture discriminative features among them and generate a compact representation of the entire image.

Vision Transformer: Self-attention-based architectures, especially Transformer [Vaswani et al., 2017] has shown major success for various Natural Language Processing (NLP) [Brown et al., 2020] as well as for Computer Vision tasks [Alamri et al., 2021, Dosovitskiy et al., 2021]; the reader is referred to [Khan et al., 2021] for further reading on Vision Transformer based literature. Specifically, CaiT [Touvron et al., 2021] introduces deeper transformer networks, and Swin Transformer [Liu et al., 2021] proposes a hierarchical Transformer, where the representation is computed using self-attention via shifted windows. In addition, TNT [Han et al., 2021] proposes transformer-backbone method modelling not only the patch-level features but also the pixel-level representations. CrossViT [Chen et al., 2021] shows how dual-branch Transformer combining different sized image patches produce stronger image features. Since the applicability of transformer-based models is growing, we aim to expand and judge its capability for GZSL tasks; to the best of our knowledge, this is still unexplored. Therefore, different from the existing works, we employ ViT to map the visual information to the semantic space, benefiting from the great performance of multi-head self-attention to learn class-level attributes.

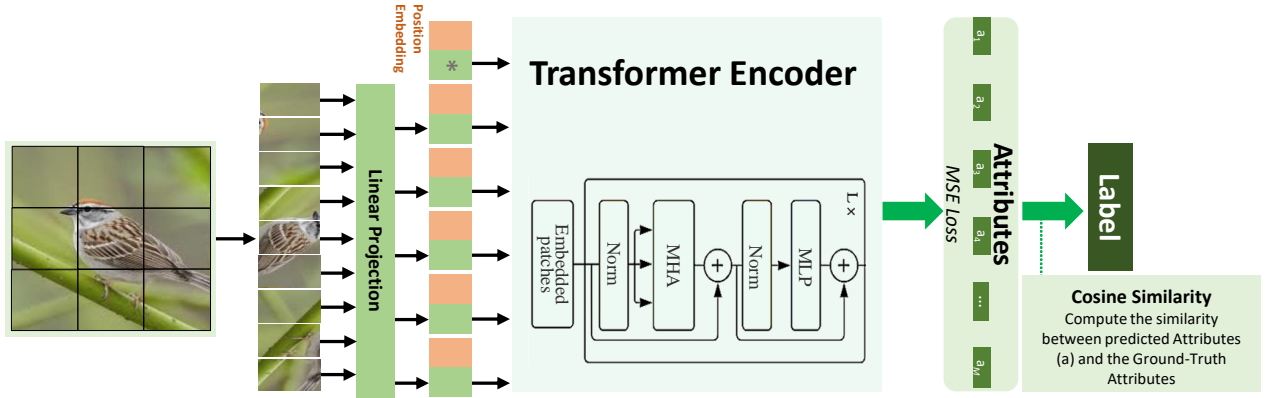


Figure 2: ViT-ZSL Architecture. An image is split into small patches fed into the Transformer encoder after attaching positional embeddings. During the training the output of the encoder is compared with the semantic information of the corresponding image via MSE loss. At inference the encoder output is used to search for the nearest class label.

3 Vision Transformer for Zero-shot Learning (ViT-ZSL)

We follow the inductive approach for training our model, i.e. during training, the model only has access to the images and corresponding image/object attributes from the *seen* classes $\mathcal{S} = \{\mathbf{x}, \mathbf{y} | \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}^s\}$, where \mathbf{x} is an RGB image and \mathbf{y} is the class-level attribute vector annotated with M different attributes, as provided with the dataset. As depicted in Figure 2, a 224×224 image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ with resolution $H \times W$ and C channels is fed into the model. The model follows ViT [Dosovitskiy et al., 2021] as closely as possible; hence the image is divided into a sequence of N patches denoted as $\mathbf{x}_p \in \mathbb{R}^{N \times (P^2 \cdot C)}$, where $N = \frac{H \cdot W}{P^2}$. Each patch with a resolution of $P \times P$ is encoded into a patch embedding by a trainable 2D convolution layer (i.e., Conv2d with kernel size=(16, 16) and stride=(16, 16)). Position embeddings are then attached to the patch embeddings to preserve the relative positional information of the order of the sequence due to the lack of recurrence in the Transformer. An extra learnable classification token ($\mathbf{z}_0^0 = \mathbf{x}_{\text{class}}$) is appended at the beginning of the sequence to encode the global image representation. Patch embeddings (\mathbf{z}) are then projected through a linear projection \mathbf{E} to D dimension (i.e., $D = 1024$) as in Eq. 1. Embeddings are then passed to the Transformer encoder, which consists of Multi-Head Attention (MHA) (Eq. 2) and MLP blocks (Eq. 3). Before every block, a layer normalisation (Norm) is employed, and residual connections are also applied after every block. Image representation ($\hat{\mathbf{y}}$) is produced as in Eq. 4.

$$\mathbf{z}_0 = [\mathbf{x}_{\text{class}}; \mathbf{x}_p^1 \mathbf{E}; \mathbf{x}_p^2 \mathbf{E}; \mathbf{x}_p^3 \mathbf{E}; \dots; \mathbf{x}_p^N \mathbf{E}] + \mathbf{E}_{\text{pos}}, \quad \mathbf{E} \in \mathbb{R}^{(P^2 \cdot C) \times D}, \mathbf{E}_{\text{pos}} \in \mathbb{R}^{(N+1) \times D} \quad (1)$$

$$\mathbf{z}'_\ell = \text{MHA}(\text{Norm}(\mathbf{z}_{\ell-1})) + \mathbf{z}_{\ell-1}, \quad \ell = 1 \dots L \quad (L = 24) \quad (2)$$

$$\mathbf{z}_\ell = \text{MLP}(\text{Norm}(\mathbf{z}'_\ell)) + \mathbf{z}'_\ell, \quad \ell = 1 \dots L \quad (3)$$

$$\hat{\mathbf{y}} = \text{Norm}(\mathbf{z}_L^0) \quad (4)$$

In terms of MHA, self-attention is performed for every patch in the sequence of the patch embeddings independently; thus, attention works simultaneously for all the patches, leading to multi-head self-attention. Three vectors, namely Query (Q), Key (K) and Value (V), are created by multiplying the encoder's input (i.e., patch embeddings) by three weight matrices (i.e., W^Q , W^K and W^V) trained during the training process to compute the self-attention. The Q and K vectors undergo a dot-product to output a scoring matrix representing how much a patch embedding has to attend to every other embedding; the higher the score is, the more attention is considered. The score matrix is then scaled down and passed into a softmax to convert the scores into probabilities, which are then multiplied by the V vectors, as in Eq. 5, where d_k is the dimension of the K vectors. Since the multi-attention mechanism is employed, self-attention matrices are then concatenated and

fed into a linear layer and passed to the regression head.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (5)$$

We argue that self-attention allows our model to attend to image regions that can be semantically relevant for classification and learns the visual features across the entire image. Since the standard ViT has one classification head implemented by an MLP, it has been edited to meet our model objective: to predict M number of attributes (i.e., depending on the datasets used). The motivation behind this is that the network is assumed to learn the notion of classes to predict attributes. For the objective function, we employed the Mean Squared Error (MSE) loss, as the continuous attributes are used as in Eq. 6, where y_i is the observed attributes, and \hat{y}_i is the predicted ones.

$$\mathcal{L}_{\text{MSE}} = \frac{1}{M} \sum_{i=1}^M (y_i - \hat{y}_i)^2 \quad (6)$$

During testing, instead of applying the extensively used dot product as in [Xu et al., 2020], we consider the cosine similarity as in [Gidaris and Komodakis, 2018] to predict class labels. The cosine similarity between the predicted attributes and every class embedding is measured. The output of the similarity measure is then used to determine the class label of the test images.

4 Experiments

Implementation Details: All images used in training and testing are adapted from the ZSL datasets mentioned below and sized 224×224 without any data augmentation. We employ the Large variant of ViT (ViT-L) [Dosovitskiy et al., 2021], with input patch size 16×16 , 1024 hidden dimension, 24 layers, 16 heads on each layer, and 24 series encoder. There are 307M parameters in total in this architecture. ViT-L is then fine-tuned using Adam optimiser with a fixed learning rate of 10^{-4} and a batch size of 64. All methods are implemented in PyTorch¹ on an NVIDIA RTX 3090 GPU, Xeon processor, and a memory sized 32GB.

Datasets: We have conducted our experiments on three popular ZSL datasets: AWA2, CUB, and SUN, whose details are presented in Table 1. The main aim of this experimentation is to validate our proposed method, ViT-ZSL, demonstrate its effectiveness and compare it with the existing state-of-the-arts. Among these datasets, AWA2 [Xian et al., 2017] consists of 37,322 images of 50 categories (40 seen + 10 unseen). Each category contains 85 binary as well as continuous class attributes. CUB [Wah et al., 2011] contains 11,788 images forming 200 different types of birds, among them 150 classes are considered as seen, and the other 50 as unseen, which is split by [Akata et al., 2016]. Together with images CUB dataset also contains 312 attributes describing birds. Finally, SUN [Patterson and Hays, 2012] has the largest number of classes among others. It consists of 717 types of scene, divided into 645 seen and 72 unseen classes. The SUN dataset contains 14,340 images with 102 annotated attributes.

Table 1: Dataset statistics in terms of granularity, number of classes (seen + unseen classes) as shown within parenthesis, number of attributes and number of images.

Datasets	Granularity	# Classes (S + U)	# Attributes	# Images
AWA2 [Xian et al., 2017]	coarse	50 (40 + 10)	85	37,322
CUB [Wah et al., 2011]	fine	200 (150 + 50)	102	11,788
SUN [Patterson and Hays, 2012]	fine	717 (645 + 72)	312	14,340

Evaluation: In this work, we train our ViT-ZSL model following the inductive approach [Wang et al., 2019]. Following [Xian et al., 2019a], we measure the top-1 accuracy for both seen as well as unseen classes. To

¹Our code is available at: <https://github.com/FaisalAlamri0/ViT-ZSL>

capture the trade-off between both sets of classes performance, we use the harmonic mean, which is the primary evaluation criterion for our model. Following the recent papers (e.g., [Xu et al., 2020], [Chao et al., 2016]), we apply Calibrated Stacking [Chao et al., 2016] to evaluate the considered methods under GZSL setting, where the calibration factor γ is dataset dependant and decided based on a validation set.

Quantitative Results: We consider the AWA2, CUB and SUN datasets to show the performance of our proposed model and compare the performance with related arts. Table 2 shows the quantitative comparison between the proposed model and various other GZSL models. The performance of each model is shown in terms of Seen (S) and Unseen (U) classes and their harmonic mean (H).

Table 2: Generalised zero-shot classification performance on AWA2, CUB and SUN

Models	AWA2			CUB			SUN		
	S	U	H	S	U	H	S	U	H
DAP [Lampert et al., 2009]	84.7	0.0	0.0	67.9	1.7	3.3	25.1	4.2	7.2
IAP [Lampert et al., 2009]	87.6	0.9	1.8	72.8	0.2	0.4	37.8	1.0	1.8
DeViSE [Frome et al., 2013]	74.7	17.1	27.8	53.0	23.8	32.8	30.5	14.7	19.8
ConSE [Norouzi et al., 2014]	90.6	0.5	1.0	72.2	1.6	3.1	39.9	6.8	11.6
SSE [Zhang and Saligrama, 2015]	82.5	8.1	14.8	46.9	8.5	14.4	36.4	2.1	4.0
SJE [Akata et al., 2015]	73.9	8.0	14.4	59.2	23.5	33.6	30.5	14.7	19.8
ESZSL [Romera-Paredes and Torr, 2015]	77.8	5.9	11.0	63.8	12.6	21.0	27.9	11.0	15.8
LATEM [Xian et al., 2016]	77.3	11.5	20.0	57.3	15.2	24.0	28.8	14.7	19.5
ALE [Akata et al., 2016]	81.8	14.0	23.9	62.8	23.7	34.4	33.1	21.8	26.3
SAE [Kodirov et al., 2017]	82.2	1.1	2.2	54.0	7.8	13.6	18.0	8.8	11.8
AREN [Xie et al., 2019]	92.9	15.6	26.7	78.7	38.9	52.1	38.8	19.0	25.5
SGMA [Zhu et al., 2019]	87.1	37.6	52.5	71.3	36.7	48.5	-	-	-
APN [Xu et al., 2020]	78.0	56.5	65.5	69.3	65.3	67.2	34.0	41.1	37.6
*GAZSL [Zhu et al., 2018]	86.5	19.2	31.4	60.6	23.9	34.3	34.5	21.7	26.7
*f-CLSWGAN [Xian et al., 2018]	64.4	57.9	59.6	57.7	43.7	49.7	36.6	42.6	39.4
Our model (ViT-ZSL)	90.0	51.9	65.8	75.2	67.3	71.0	55.3	44.5	49.3

S, U, H denote Seen classes (\mathcal{Y}^s), Unseen classes (\mathcal{Y}^u), and the Harmonic mean, respectively. For each scenario, the best is in **red** and the second-best is in **blue**. * indicates generative representation learning methods.

DAP and IAP [Lampert et al., 2009] are some of the earliest works in ZSL, which perform poorly compared to other models. This is due to the assumptions claimed in these approaches regarding attributes dependency. In real-world animals with attributes ‘terrestrial’ and ‘farm’ are dependent but are assumed independent by such models, which are noted as incorrect by [Akata et al., 2016]. Our model ViT-ZSL does not assume this, but rather it considers the correlation between attributes, which self-attention helps to achieve by considering both positional and contextual information of the entire sequence of patches. DeVISE [Frome et al., 2013] and ConSE [Norouzi et al., 2014] learn a linear mapping between images and their semantic embedding space. They both make use of the same text model trained on 5.4B words from Wikipedia to construct 500-dimensional word embedding vectors. Both use the same baseline model, but DeVISE replaces the last layer (i.e., softmax layer) with a linear transformation layer. In contrast, ConSE keeps it and computes the predictions via a convex combination of the class label embedding vectors. ConSE, as presented in Table 2 outperforms DeVISE, but DeVISE is generally performing better on the unseen classes. Similarly, SJE [Akata et al., 2015] learns a bilinear compatibility function using the structural SVM objective function to maximise the compatibility between image and class embeddings. ESZSL [Romera-Paredes and Torr, 2015] uses the square loss to learn bilinear compatibility. Although ESZSL is claimed to be easy to implement, its performance, in particular for GZSL, is poor. ALE [Akata et al., 2016], which belongs to the bilinear compatibility approach group, performs better than most of its group member. LATEM [Xian et al., 2016], instead of learning a single bilinear map, extends the bilinear compatibility of SJE [Akata et al., 2015] as to be an image-class pairwise linear by learning multiple linear mappings. It performs better than SJE on unseen classes but with a lower harmonic

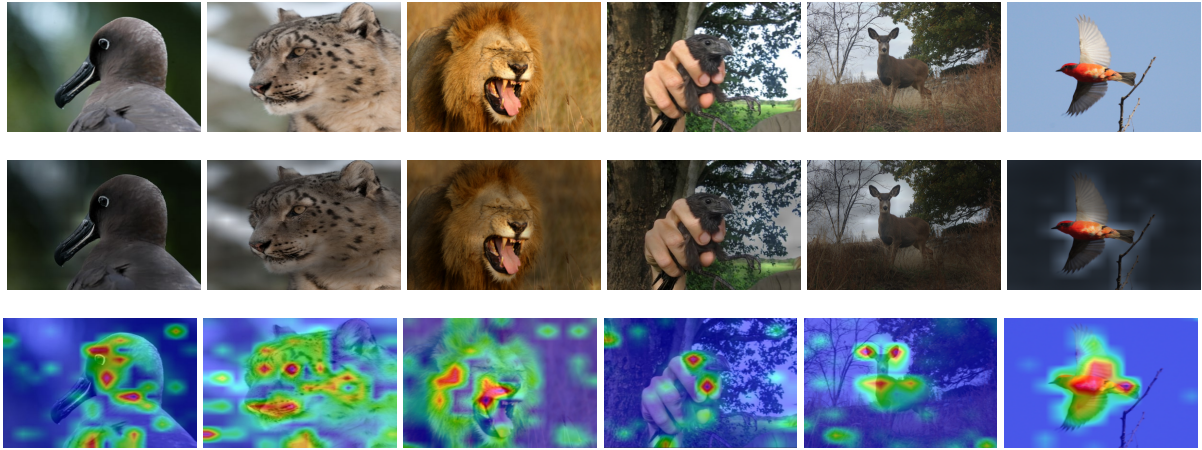


Figure 3: Representative examples of attention. First row: Original images, Middle: Attention maps, and last: Attention fusions. From left to right, ViT-ZSL is able to focus on object-level attributes and learn objects discriminative features when objects are partly captured (first three columns images), occluded (fourth column images) or fully presented (last two columns images).

mean due to its poor performance on seen classes. Generative ZSL models such as GAZSL [Zhu et al., 2018], and f-CLSWGAN [Xian et al., 2018] are seen to reduce the effect of the bias problem due to the inclusion of synthesised features for the unseen classes. However, this does not apply to our method, as no synthesised features are used in our case; instead, solely the features extracted from seen classes are used during training. AREN [Xie et al., 2019], SGMA [Zhu et al., 2019] and APN [Xu et al., 2020] are non-generative ZSL models focusing on object region localisation using image attention. They are the most relevant works to ours as attention mechanism is included in these models architecture. However, they consist of two branches in their models, where the first learns local discriminative visual features and the second captures the image’s global context. In contrast, our model uses only one compact network, where the input is the image patches so that the global and local discriminative features can be learned using the multi-head self-attention mechanism.

Our model ViT-ZSL, as shown in Table 2, achieves the best harmonic mean on AWA2. It also performs as the third best on both seen and unseen classes. Compared with the other models, it scores 90.02%, where the highest is the highest is AREN with 92.9% accuracy. As the comparison illustrated follows the GZSL setting using the harmonic mean as the primary evaluation metric for GZSL models, ViT-ZSL outperforms all state-of-the-art models. In terms of the CUB dataset, our method achieves the second-highest accuracy for seen classes, but the highest for unseen. In addition, our ViT-ZSL obtains the best harmonic mean score among all the reported approaches. On SUN datasets, which has the most significant number of object classes among other datasets, our model performs as the best for both seen and unseen classes, achieving a harmonic mean of 47.9%, the highest compared to all other models.

Attention Maps: In Figure 3, we show how our model attends to image regions semantically relevant to the object class. For example, in the images of the first three columns, the entire objects’ shapes are absent (i.e., only the top part is captured), and in the image in the fourth column, the groove-billed ani bird is impeded by a human hand. Although these images suffer from occlusion, our model accurately attends to the objects in the image. Thus, we believe that ViT-ZSL definitely benefits from the attention mechanism, which is also involved in the human recognition system. Clearly, we can say that our method has learned to map the relevance of local regions to representations in the semantic space, where it makes predictions on the visible attribute-based regions. Similarly, in the last two columns images of Figure 3, it can be noticed how the model pays more attention to some object-level attributes (i.e., *Deer*: forest, agility, furry etc., and *Vermilion Flycatcher*: solid and red breast, perching-like shape, notched tail). It can also be noticed that the model focuses on the context of the object, as in the second column images. This can be due to the guidance of some attributes (i.e., forest, jungle, ground and tree) which are associated with *leopard* class. However, as shown in the first column, the model did not pay much attention to the bird’s beak compared to the head and the rest of the body, which

needs to be investigated further and building an explainable model as in [Xian et al., 2019b] could be a way to accomplish this.

5 Conclusion

In this paper, we proposed a Vision Transformer-based Zero-Shot Learning (ViT-ZSL) model that specifically exploits the multi-head self-attention mechanism for relating visual and semantic attributes. Our qualitative results showed that the attention mechanism involved in our model focuses on the most relevant image regions related to the object class to predict the semantic information, which is used to find out the class label during inference. Our results on the GZSL task, including the highest harmonic mean scores on the AWA2, CUB and SUN datasets, illustrate the effectiveness of our proposed method.

Although our method achieves very encouraging results for the GZSL task on three publicly available benchmarks, the bias problem towards seen classes remains a challenge, which will be addressed in future work. Training the model in a transductive setting, where visual information for unseen classes could be included, is a direction to be examined.

Acknowledgement

This work was supported by the Defence Science and Technology Laboratory and the Alan Turing Institute. The TITAN Xp and TITAN V used for this research were donated by the NVIDIA Corporation.

References

- [Akata et al., 2016] Akata, Z., Perronnin, F., Harchaoui, Z., and Schmid, C. (2016). Label-embedding for image classification. *IEEE TPAMI*.
- [Akata et al., 2015] Akata, Z., Reed, S. E., Walter, D., Lee, H., and Schiele, B. (2015). Evaluation of output embeddings for fine-grained image classification. In *CVPR*.
- [Alamri et al., 2021] Alamri, F., Kalkan, S., and Pugeault, N. (2021). Transformer-encoder detector module: Using context to improve robustness to adversarial attacks on object detection. In *ICPR*.
- [Brown et al., 2020] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. In *NeurIPS*.
- [Chao et al., 2016] Chao, W.-L., Changpinyo, S., Gong, B., and Sha, F. (2016). An empirical study and analysis of generalized zero-shot learning for object recognition in the wild. In *ECCV*.
- [Chen et al., 2021] Chen, C.-F., Fan, Q., and Panda, R. (2021). Crossvit: Cross-attention multi-scale vision transformer for image classification. *arXiv*.
- [Dosovitskiy et al., 2021] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*.
- [Federici et al., 2020] Federici, M., Dutta, A., Forré, P., Kushman, N., and Akata, Z. (2020). Learning Robust Representations via Multi-View Information Bottleneck. In *ICLR*.
- [Frome et al., 2013] Frome, A., Corrado, G. S., Shlens, J., Bengio, S., Dean, J., Ranzato, M. A., and Mikolov, T. (2013). Devise: A deep visual-semantic embedding model. In *NIPS*.
- [Gidaris and Komodakis, 2018] Gidaris, S. and Komodakis, N. (2018). Dynamic few-shot visual learning without forgetting. In *CVPR*.
- [Han et al., 2021] Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., and Wang, Y. (2021). Transformer in transformer. *arXiv*.
- [Khan et al., 2021] Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F., and Shah, M. (2021). Transformers in vision: A survey. *arXiv*.

- [Kodirov et al., 2017] Kodirov, E., Xiang, T., and Gong, S. (2017). Semantic autoencoder for zero-shot learning. In *CVPR*.
- [Lampert et al., 2009] Lampert, C. H., Nickisch, H., and Harmeling, S. (2009). Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*.
- [Liu et al., 2021] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv*.
- [Mikolov et al., 2013] Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *NIPS*.
- [Norouzi et al., 2014] Norouzi, M., Mikolov, T., Bengio, S., Singer, Y., Shlens, J., Frome, A., Corrado, G., and Dean, J. (2014). Zero-shot learning by convex combination of semantic embeddings. In *ICLR*.
- [Patterson and Hays, 2012] Patterson, G. and Hays, J. (2012). Sun attribute database: Discovering, annotating, and recognizing scene attributes. In *CVPR*.
- [Pennington et al., 2014] Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *EMNLP*.
- [Ren et al., 2015] Ren, S., He, K., Girshick, R. B., and Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE TPAMI*.
- [Romera-Paredes and Torr, 2015] Romera-Paredes, B. and Torr, P. (2015). An embarrassingly simple approach to zero-shot learning. In *ICML*.
- [Schönfeld et al., 2019] Schönfeld, E., Ebrahimi, S., Sinha, S., Darrell, T., and Akata, Z. (2019). Generalized zero- and few-shot learning via aligned variational autoencoders. *CVPR*.
- [Touvron et al., 2021] Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., and Jégou, H. (2021). Going deeper with image transformers. *arXiv*.
- [Vaswani et al., 2017] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. In *NIPS*.
- [Wah et al., 2011] Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. (2011). The Caltech-UCSD Birds-200-2011 Dataset. Technical report, California Institute of Technology.
- [Wang et al., 2019] Wang, W., Zheng, V., Yu, H., and Miao, C. (2019). A survey of zero-shot learning. *ACM-TIST*.
- [Xian et al., 2016] Xian, Y., Akata, Z., Sharma, G., Nguyen, Q., Hein, M., and Schiele, B. (2016). Latent embeddings for zero-shot classification. In *CVPR*.
- [Xian et al., 2019a] Xian, Y., Lampert, C. H., Schiele, B., and Akata, Z. (2019a). Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE TPAMI*.
- [Xian et al., 2018] Xian, Y., Lorenz, T., Schiele, B., and Akata, Z. (2018). Feature generating networks for zero-shot learning. In *CVPR*.
- [Xian et al., 2017] Xian, Y., Schiele, B., and Akata, Z. (2017). Zero-shot learning - the good, the bad and the ugly. In *CVPR*.
- [Xian et al., 2019b] Xian, Y., Sharma, S., Schiele, B., and Akata, Z. (2019b). F-vaegan-d2: A feature generating framework for any-shot learning. In *CVPR*.
- [Xie et al., 2019] 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*.
- [Xu et al., 2020] Xu, W., Xian, Y., Wang, J., Schiele, B., and Akata, Z. (2020). Attribute prototype network for zero-shot learning. In *NIPS*.
- [Yu et al., 2018] Yu, y., Ji, Z., Fu, Y., Guo, J., Pang, Y., and Zhang, Z. M. (2018). Stacked semantics-guided attention model for fine-grained zero-shot learning. In *NeurIPS*.
- [Zhang and Saligrama, 2015] Zhang, Z. and Saligrama, V. (2015). Zero-shot learning via semantic similarity embedding. In *ICCV*.
- [Zhu et al., 2018] Zhu, Y., Elhoseiny, M., Liu, B., Peng, X., and Elgammal, A. (2018). Imagine it for me: Generative adversarial approach for zero-shot learning from noisy texts. In *CVPR*.
- [Zhu et al., 2019] Zhu, Y., Xie, J., Tang, Z., Peng, X., and Elgammal, A. (2019). Semantic-guided multi-attention localization for zero-shot learning. In *NIPS*.