

# Epsilon: Exploring Comprehensive Visual-Semantic Projection for Multi-Label Zero-Shot Learning

Ziming Liu<sup>1</sup>, Jingcai Guo<sup>1\*</sup>, Song Guo<sup>2</sup>, Xiaocheng Lu<sup>2</sup>

<sup>1</sup>Department of Computing, The Hong Kong Polytechnic University

<sup>2</sup>Department of Computer Science and Engineering, The Hong Kong University of Science and Technology  
ziming.liu@connect.polyu.hk, jc-jingcai.guo@polyu.edu.hk, songguo@cse.ust.hk, xlubl@connect.ust.hk

## Abstract

This paper investigates a challenging problem of zero-shot learning in the multi-label scenario (MLZSL), wherein the model is trained to recognize multiple unseen classes within a sample (e.g., an image) based on seen classes and auxiliary knowledge, e.g., semantic information. Existing methods usually resort to analyzing the relationship of various seen classes residing in a sample from the dimension of spatial or semantic characteristics and transferring the learned model to unseen ones. However, they neglect the integrity of local and global features. Although the use of the attention structure will accurately locate local features, especially objects, it will significantly lose its integrity, and the relationship between classes will also be affected. Rough processing of global features will also directly affect comprehensiveness. This neglect will make the model lose its grasp of the main components of the image. Relying only on the local existence of seen classes during the inference stage introduces unavoidable bias. In this paper, we propose a novel and comprehensive visual-semantic framework for MLZSL, dubbed *Epsilon*, to fully make use of such properties and enable a more accurate and robust visual-semantic projection. In terms of spatial information, we achieve effective refinement by group aggregating image features into several semantic prompts. It can aggregate semantic information rather than class information, preserving the correlation between semantics. In terms of global semantics, we use global forward propagation to collect as much information as possible to ensure that semantics are not omitted. Experiments on large-scale MLZSL benchmark datasets *NUS-Wide* and *Open-Images-v4* demonstrate that the proposed *Epsilon* outperforms other state-of-the-art methods with large margins.

## Introduction

With the continuous advances in science and technology, the ability and ways human beings obtain information have greatly improved. In recent years, deep learning classification networks (Simonyan and Zisserman 2014; Szegedy et al. 2015) driven by computational intelligence have sprung up and made rapid progress. Image classification datasets designed for single labels in the past are less challenging for deep learning models. Since an arbitrary image

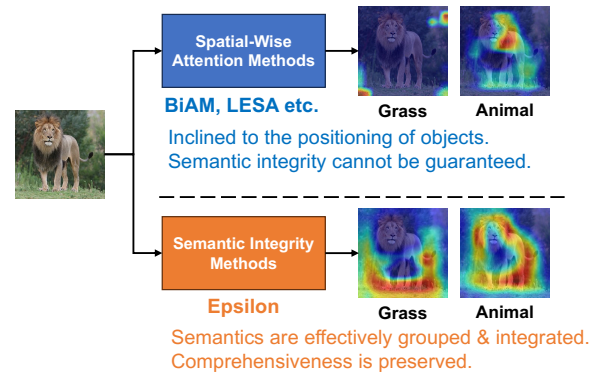


Figure 1: Comparison of attention between our proposed *Epsilon* and traditional spatial attention-based models. We can see that the proposed model is stronger than the contrasting *BiAM* and *LESA* in terms of completeness of class semantics. (zoom in for a better view)

may contain multiple objects, the task of multi-label image classification (Gong et al. 2013; Yu et al. 2014; Weston, Bengio, and Usunier 2011) has come into view because the task is challenging, and it is closer to the real life of a human. With the continuous improvement of deep neural networks, multi-label image classification models based on deep learning have emerged (Wang et al. 2016; Durand, Mehrasa, and Mori 2019; Gao and Zhou 2021; Feng, An, and He 2019). Their research focuses on how to effectively process the context information in the image, or build a graph model or other topological models to handle complex dependencies between labels (Chen et al. 2019). In recent years, researchers have realized that the attention mechanism has been introduced to allow the model to focus more selectively on information-rich regions and has achieved remarkable results (Cheng et al. 2021).

However, the above multi-label image classification methods all have a common problem, that is, these models cannot perceive categories that were not present during training. Therefore, zero-shot learning (ZSL) has become a solution to this problem (Akata et al. 2016; Deutsch et al. 2017; Frome et al. 2013). ZSL is essentially a multi-modal transfer learning, that is, an algorithm that predicts unseen classes during testing through the learning of seen classes. There-

\*Corresponding Author.

fore, it is different from the tasks and challenges faced by traditional supervised learning. In fact, existing ZSL models have achieved substantial success in single object prediction (Kodirov, Xiang, and Gong 2017; Xian, Schiele, and Akata 2017; Guo and Guo 2020). Compared with traditional ZSL, generalized ZSL is closer to the actual application scenario. Generalized ZSL means that not only unseen classes need to be predicted, but also all seen classes should be detected during testing. However, in practical problems, there are often multiple seen objects and unseen objects in a picture. At this time, the traditional ZSL model will cause serious performance loss and even cannot work effectively. Therefore, multi-label zero-shot learning (MLZSL) has entered the research field as a more complex and practical research direction.

There have been some works focusing on MLZSL tasks and obtained some achievements with decisive progress in recent years (Norouzi et al. 2014; Mensink, Gavves, and Snoek 2014; Zhang, Gong, and Shah 2016). Among them, Zhang *et al.* (Zhang, Gong, and Shah 2016) and Ben *et al.* (Ben-Cohen et al. 2021) calculate the principal direction of its feature vectors for each test image and then sorts the similarity of labels that exist near that principal direction. *LESA* (Huynh and Elhamifar 2020) and *BiAM* (Narayan et al. 2021), on the other hand, performs multi-label prediction tasks from another perspective, which strengthens the feature response of each class through attention mechanisms. (Liu et al. 2023) innovatively handles the semantic relationship between seen and unseen classes from the perspective of channel response, but this approach is not precise enough for the processing of spatial features.

In response to the problems encountered in the MLZSL method mentioned above, the first problem is that previous methods did not effectively utilize the relationship between local and global features of images. For example, *SDL* (Ben-Cohen et al. 2021) and *Fast0Tag* (Zhang, Gong, and Shah 2016) only focus on the principal vector of the image, ignoring the supplement of local features to global information. However, *BiAM* (Narayan et al. 2021) simply concatenates global and local features and then performs the classification part. *LESA* (Huynh and Elhamifar 2020) uses the attention sharing of local features to replace the global features. In the multi-label environment, global features are the most important for determining the main components in an image. This ensures that the semantic integrity is not destroyed. As shown in Figure 1, comparing the attention of the two classes of ‘grass’ and ‘animal’, it is not difficult to see that the overall semantics of the spatial attention-based method are not comprehensive. The consequence is that it will affect the model’s learning of the seen class and its ability to transfer to the unseen class. Similar problems occur extremely frequently, especially for semantic extraction of scene classes. Therefore, we urgently need to locally group and integrate the semantics within the image and expand the diversity of global features to make the semantic information corresponding to the image more comprehensive, thereby helping the model to better perform transfer learning. Based on this, we proposed Exploring Comprehensive Visual-Semantic Projection for Multi-Label

Zero-Shot Learning (*Epsilon*) to try our best to find semantic descriptions corresponding to all classes in the picture.

The main contributions of this paper can be summarized as the following three-fold:

1. We propose to use visual prompt learning to group and aggregate features, which solves the loss of details caused by traditional methods using spatial attention and greatly ensures semantic integrity.
2. We use the Global Forward Propagation (GFP) module to greatly enrich the diversity of global features and improve the global information richness.
3. Both modules work together on the fused features. Extensive experiments on *NUS-Wide* and *Open-Images-V4* datasets demonstrate the effectiveness of our method against other state-of-the-art MLZSL models.

## Related Work

MLZSL presents a more complex challenge than single-label learning due to the unpredictability of the number of labels assigned to each image, thus demanding the model’s capacity to weigh multiple unseen labels simultaneously. Several methods have emerged to address the challenge. Norouzi *et al.* (Norouzi et al. 2014) developed a model for multi-label zero-shot learning by explicitly partitioning the image and semantic embedding spaces and using a convex combination of label embedding vectors to map the image. Zhang *et al.* (Zhang, Gong, and Shah 2016) proposed a more general and fast model based on word vectors of ranked relevant labels. Lee *et al.* (Lee et al. 2018) posited that knowledge graphs offer a framework for linking different labels in multi-label environments. He is also the first model to introduce graph structure into multi-label zero-shot learning. Attention-based methods such as *LESA* (Huynh and Elhamifar 2020) and Narayan *et al.* (Narayan et al. 2021) have shown promising results. Their approaches leverage the attention-sharing mechanism and bi-layer attention module to focus on key areas of different labels and global context information, respectively. Compared with past methods, they pay more attention to the mining of local semantics and the sharing of features between classes. Although Ben *et al.* (Ben-Cohen et al. 2021) utilized the diversity of semantics and embedding matrices to enhance the ability of multi-label zero-shot learning, such methods require analysis in real-world scenarios. Nevertheless, these approaches fall short of exploring the feature channels’ response to various classes and only stay at the two-dimensional ( $H \times W$ ) level. Liu *et al.* (Liu et al. 2023), starting from the perspective that excessive mining of spatial features will cause loss to image semantics, creatively proposed to find inter-class relationships from the perspective of channel response. Channel responses can also be encoded into effective semantic information in images, which greatly improves model efficiency. Gupta *et al.* (Gupta et al. 2023) creatively extended the GAN model of single-label zero-shot learning to a multi-label environment. And the multi-label information is fused at the attribute-level and feature-level respectively, and finally summarized at the cross-level.

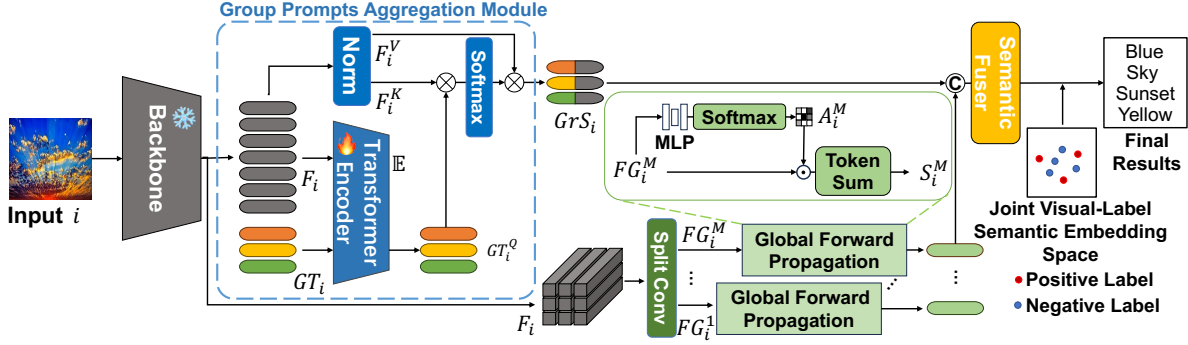


Figure 2: Pipeline of *Epsilon*. The feature representation of the image is first obtained through a pre-trained backbone network that is frozen and cannot be updated. The image features are then applied to the Group Prompts Aggregation Module (GPA Module), which represents local semantic generation, and the Global Forward Propagation Module (GFP Module), which represents global diversity semantic generation. Finally, the output of the two modules is integrated to obtain the complete semantics. (zoom in for a better view)

## Methods

### Preliminary Problem Setting

It is necessary to clearly define the purpose of the MLZSL. Let the batch size be  $n$ , for the images input into the model  $\{(I_1, Y_1), \dots, (I_i, Y_i), \dots, (I_n, Y_n)\}$ , where  $I_i$  is the  $i$ -th image in the input training set, and the corresponding  $Y_i$  represents the label of the input image  $i$ . These labels in the training set are also called ‘seen labels’. Like ZSL, MLZSL does not overlap the labels of the training and testing sets in terms of label distribution. Let’s define the set of all labels in the dataset as  $\mathcal{C}$ , the set of seen labels as  $\mathcal{C}_s$ , and the set of unseen labels as  $\mathcal{C}_u$ . The distribution relationship of labels in the dataset can be described as  $\mathcal{C} = \mathcal{C}_s \cup \mathcal{C}_u$ .  $\mathcal{C}_s$  is mainly used for training sets, while  $\mathcal{C}_u$  mainly appears during the testing process. During the MLZSL testing process, for the input image  $I_u$ , the output prediction result  $y_u$  should be  $y_u \subset \mathcal{C}_u$ . In the generalized MLZSL task, the output prediction result  $y_u$  must include ‘seen labels’, which is  $y_u \subset \mathcal{C}$ .

### Group Prompts Aggregation Module

In order to obtain finer-grained spatial features, an effective semantic grouping of spatial features is performed. In this section, we introduce the concept of feature grouping. The specific structure of the model is shown in Figure 2. First, for the input image  $I_i$ , we obtain the corresponding feature  $F_i$  of the image through a pre-trained backbone network (in this paper, we use the pre-trained ViT-B/16 model (Dosovitskiy et al. 2020) as an example.). Subsequently, we designed several group prompts that can be updated, with the number of group prompts being  $M$ . Here,  $M$  is one of the hyper-parameters of the model, and we will discuss it in the experimental session. In order to better group and aggregate image features, we send the image features  $F_i \in \mathbb{R}^{B \times N \times D}$  and  $GT_i \in \mathbb{R}^{B \times M \times D}$  representing group prompts to the encoder of Transformer model for aggregation, where  $N$  is the number of image tokens,  $D$  is the feature dimension of the token. It should be noted here that the Transformer’s Encoder can be updated. This is because we need the model itself to learn how to integrate features rather than just relying

on updating the prompts themselves. The above processes are expressed as:

$$\mathcal{T}_i = [GT_i; F_i], \quad (1)$$

$$GT_i^Q = \mathbb{E}(\mathcal{T}_i)[: M, :], \quad (2)$$

Among them,  $[: M, :]$  represents the first  $M$  prompts output, and  $\mathbb{E}$  represents the encoder of the Transformer model.  $GT_i^Q \in \mathbb{R}^{B \times M \times D}$  is the aggregated image feature we obtained. However, the above operations can only roughly group features.

In order to obtain more refined group features for subsequent encoding of semantic information, we recombine these updateable group prompts  $GT_i^Q$  with the original features  $F_i$ . Here,  $GT_i^Q$  serves as query and  $F_i$  serves as key and value. This step further refines the semantic information of each group in order to complete visual-semantic projection locally.

$$GrS_i = softmax(F_i^K \cdot GT_i^Q) \cdot F_i^V \quad (3)$$

Among them,  $GrS_i \in \mathbb{R}^{B \times M \times D}$  is the group semantic vector, which is a clustering of local semantic information.

### Global Forward Propagation Module

In existing methods, the extraction of global features is rough. Most methods use the top-level features and design some channel (or global) attention structures to achieve this. The above approach ignores the diversity of global information. In fact, the global information itself should have sufficient information abundance to reflect its comprehensiveness. Therefore, we designed the Global Forward Propagation Module (GFP Module) to implement an extraction strategy for richer global information. This new structural design feeds existing global features into multiple modules respectively and uses the personalized global weights obtained by the modules to enhance global information from multiple angles. The specific structure is shown in Figure 2. We first

encode the feature  $F \in \mathbb{R}^{B \times N \times D}$  into  $M$  feature blocks, respectively  $FG_i^1, FG_i^2, \dots, FG_i^M \in \mathbb{R}^{B \times N \times D}$ . Then,  $M$  identical modules are applied to each feature block. The specific implementation process of the module is as follows: Taking  $FG_i^M \in \mathbb{R}^{B \times N \times D}$  as an example, we first use a multi-layer perceptron (MLP) to re-represent the features. Then, the features are fed into softmax to obtain the weight representation of each feature point in the feature map  $A_i^M \in \mathbb{R}^{B \times N \times D}$ . The reason for this operation is that we need the GFP module to improve the diversity of global features during the forward propagation process. At the same time, we must also ensure that important global information is given the necessary prominence. Therefore, this personalized weight learning and structure design meets our needs. After obtaining  $A_i^M$ , we perform a dot product with the original input feature  $FG_i^M$ . Then, all features are summed to obtain the global semantics  $S_i^M \in \mathbb{R}^{B \times D}$  corresponding to  $FG_i^M$ . The above processes are expressed as:

$$W_i^M = \mathbb{M}(FG_i^M), \quad (4)$$

$$A_i^M = \text{softmax}(W_i^M), \quad (5)$$

$$S_i^M = \text{Sum}(FG_i^M \cdot A_i^M). \quad (6)$$

Among them,  $\cdot$  means the dot product operation.  $\mathbb{M}$  means the MLP structure. After obtaining all the global semantics  $GoS_i = S_i^1, \dots, S_i^M$ , we use the concatenate operation to combine  $GoS_i$  and  $GrS_i$  and get  $GS_i \in \mathbb{R}^{B \times M \times 2D}$ . This idea is mainly referred to Ben *et al.* (Ben-Cohen et al. 2021) and Zhang *et al.* (Zhang, Gong, and Shah 2016), which both use the principal vector of the image as the semantic information. Although our GFP module has some structural similarities with the traditional channel attention model, our distribution of weights is not limited to the channel dimension. In addition, since traditional attention models tend to let the model learn a global weight distribution, this will cause the global features of different images to lose diversity. Our multi-head approach enriches this diversity.

After concatenating the semantic information, we pass a Semantic Fuser, which is a linear layer that helps change the length of the existing semantics so that it is the same as the length of the class semantics in the semantic embedding space, which facilitates the subsequent prediction process. The process is expressed as below, where  $\mathbb{C}$  refers to the concat layer, and  $\mathbb{L}$  refers to the linear layer:

$$GS_i = \mathbb{C}(GoS_i, GrS_i), \quad (7)$$

$$S_i = \mathbb{L}(GS_i), \quad (8)$$

## Loss Function

Our loss function is inspired by (Zhang, Gong, and Shah 2016; Ben-Cohen et al. 2021). For each input image  $I_i$ , the semantic group of the corresponding image is  $\mathcal{S}_i \in \mathbb{R}^{M \times d_w}$ , where  $d_w$  is the length of label embeddings. This semantic vector group is used to predict unseen classes during the testing process. Therefore, the basic calculation process of the loss function includes the judgment of class similarity. It is

necessary to try to improve the ranking of the classes that appear in the input image and reduce the similarity ranking of the classes that do not appear,

$$\tau_{jk} = \max(n_j \mathcal{S}_i) - \max(p_k \mathcal{S}_i), \quad (9)$$

where  $n_j$  is the word vector of the seen classes that do not appear, and  $p_k$  is the word vector of the seen classes that appear in the input image. In order to maximize the distance between classes that appear and those that do not, we use the maximum value for calculation. At the same time, this approach can ensure sufficient semantic diversity. Therefore, one of the main components of the loss function, namely the ranknet loss (Zhang, Gong, and Shah 2016), is shown in the following function:

$$\mathcal{L}_r = \alpha \sum_j \sum_k \log(1 + e^{\tau_{jk}}), \quad (10)$$

where  $\alpha = (|T| |\bar{T}|)^{-1}$ .  $|T|$  and  $|\bar{T}|$  denote the number of seen classes that appear or not, respectively. The hyper-parameter  $\alpha$  is used to normalize the ranknet loss.

In addition, in order to reduce the difficulty of the model in predicting hard-positive samples with a high diversity of labels, we suggest introducing a new weight like (Ben-Cohen et al. 2021). Due to the uniqueness of multi-label tasks, each image contains a different number of labels. When a large number of classes appear in the image, it is necessary to make the model sensitive to label diversity and have a tendency towards hard-positive tasks during the training process.

$$\omega = 1 + \sum_i \text{var}(Y_i). \quad (11)$$

In addition, the loss function of the model introduces the second main component, namely the regularized loss function. Its purpose is to construct correlations between input semantic vectors.

$$\mathcal{L}_{reg} = \left\| \sum_M \text{var}(\mathcal{S}_i^M) \right\|_1. \quad (12)$$

Finally, our complete loss function  $\mathcal{L}$  is as follows:

$$\mathcal{L} = \frac{1}{B} \sum_{i=1}^B (\omega \cdot (1 - \lambda) \mathcal{L}_r(\mathcal{S}_i, Y_i) + \lambda \mathcal{L}_{reg}(\mathcal{S}_i)), \quad (13)$$

where  $\lambda$  is the weight of the regularized loss function, and  $B$  is the batch size.

## Experiments

### Experimental Setup

**Datasets:** The *NUS-Wide* dataset (Chua et al. 2009) contains approximately 270,000 images and a total of 1,006 labels. Among them, 81 labels manually annotated by humans will serve as labels for ‘unseen classes’. At the same time, these labels will also serve as ‘ground-truth’ labels in the multi-label classification task. The remaining 925 labels were automatically extracted from Flickr users’ manual annotations of these images, where they will be used as labels for ‘seen

Method	ZSL							GZSL						
	mAP	K=3			K=5			mAP	K=3			K=5		
		P	R	F1	P	R	F1		P	R	F1	P	R	F1
Fast0Tag (Zhang, Gong, and Shah 2016)	15.1	22.6	36.2	27.8	18.2	48.4	26.4	3.7	18.8	8.3	11.5	15.9	11.7	13.5
Kim <i>et al.</i> (Kim, Jun, and Zhang 2018)	10.4	20.9	33.5	25.8	16.2	43.2	23.6	3.7	17.9	7.9	10.9	15.6	11.5	13.2
LESA ApC (Huynh and Elhamifar 2020)	12.9	20.0	31.9	24.6	15.7	41.9	22.9	2.6	10.4	4.6	6.4	9.1	6.7	7.7
LESA (Huynh and Elhamifar 2020)	19.4	25.7	41.1	31.6	19.7	52.5	28.7	5.6	23.6	10.4	14.4	19.8	14.6	16.8
BiAM (Narayan et al. 2021)	25.8	26.0	41.6	32.0	<b>20.2</b>	<b>53.9</b>	29.4	8.9	25.2	11.3	15.5	21.8	16.0	18.5
SDL (Ben-Cohen et al. 2021)	<b>25.9</b>	24.2	41.3	30.5	18.8	53.4	27.8	<b>12.1</b>	27.7	<b>13.9</b>	18.5	23.0	<b>19.3</b>	21.0
Gen-MLZSL (Gupta et al. 2023)	25.7	<b>26.6</b>	<b>42.8</b>	<b>32.8</b>	20.1	<b>53.6</b>	29.3	8.9	<b>30.9</b>	<b>13.6</b>	<b>18.9</b>	<b>26.0</b>	<b>19.1</b>	<b>22.0</b>
(ML) <sup>2</sup> P-Encoder (Liu et al. 2023)	29.4	-	-	<b>32.8</b>	-	-	<b>32.3</b>	10.2	-	-	16.8	-	-	19.2
Ours	<b>33.2</b>	<b>31.1</b>	<b>41.7</b>	<b>35.6</b>	<b>24.9</b>	52.7	<b>33.8</b>	<b>14.7</b>	<b>32.8</b>	13.1	<b>18.7</b>	<b>28.1</b>	18.8	<b>22.5</b>

Table 1: Comparison of different models for multi-label ZSL and GZSL tasks on the **NUS-Wide** dataset. The ‘‘P’’ refers to the top-K ‘‘Precision’’ and ‘‘R’’ refers to the top-K ‘‘Recall’’. The best results are shown in **red**, and the second are shown in **blue**.

classes’. This setting is similar with (Huynh and Elhamifar 2020; Ben-Cohen et al. 2021). Another dataset is called the *Open-Images-V4* dataset, which is much larger than the *NUS-Wide* dataset. This dataset contains approximately 9.2 million images, of which approximately 9 million are used as the training set. The training set contains a total of 7,186 labels, ensuring that each label appears at least 100 times in the training set. These will be considered as labels for ‘seen classes’. In addition, the dataset also contains 125,456 test images and 400 ‘unseen classes’ labels. These labels are derived from the other 400 most frequent labels that did not appear in the training set, which appeared at least 75 times. The setting of *Open-Images-V4* dataset is similar with (Huynh and Elhamifar 2020; Narayan et al. 2021).

**Evaluation Metrics:** In order to ensure the unbiased comparison and scientific evaluation of the metrics themselves using our proposed method and comparison method, we will use the two most commonly used evaluation metrics in MLZSL and Multi-label Classification tasks, namely the mean Average Precision (mAP) and F1-Score (Veit et al. 2017; Huynh and Elhamifar 2020). Among them, F1-Score is the harmonic mean of the precision and recall. *top-K* F1-Score is an evaluation metric used in classification tasks to measure the accuracy of the model in predicting labels. mAP is a class-wise evaluation indicator that is used to reflect the accuracy of image image retrieval from unseen labels.

**Implementation Details:** As for the selection of backbone network, we choose the pre-trained ViT-B/16 (Dosovitskiy et al. 2020) as our backbone network. Unlike the comparison method that uses multi-stage training, our method is completely end-to-end training, saving a lot of training time and making the model more concise. We choose the Adam optimizer (Kingma and Ba 2014) as the model’s optimizer, which is suitable for large-scale datasets and requires less memory. The weight decay of the Adam optimizer is set to  $4e^{-3}$ . For the experiments of all the models in the *NUS-Wide* dataset, the entire training process requires a total of 7 epochs with a batch size of 96, and the initial learning rate is set to  $1e^{-5}$ , and then decreases by  $\frac{1}{2}$  at the 4-th epoch. In the experiments of the *Open-Images-V4* dataset, the number of epochs in the training process is set to 7. This is consistent with the settings for other comparison methods. Our optimizer’s decay rate, model’s learning rate, and batch size remain the same.

**Baselines:** In terms of baseline methods, we need to com-

pare our proposed new model with all state-of-the-art MLZSL models based on deep neural networks in recent years. The baseline method covers all research directions, including but not limited to attention mechanism, principal direction of feature vectors, etc. These comparison methods include: *Fast0Tag* (Zhang, Gong, and Shah 2016), *Kim et al.* (Kim, Jun, and Zhang 2018), *LESA Attention per Cluster* (Huynh and Elhamifar 2020), *LESA* (Huynh and Elhamifar 2020), *BiAM* (Narayan et al. 2021), *SDL* (Ben-Cohen et al. 2021), Gen-MLZSL (Gupta et al. 2023), and *(ML)<sup>2</sup>P-Encoder* (Liu et al. 2023).

### MLZSL Performance for *Epsilon*

We show the performance of both ZSL and GZSL tasks in Table 1. Firstly, for *LESA* (Huynh and Elhamifar 2020) and *BiAM* (Narayan et al. 2021), as two models with the help of spatial attention mechanism, they are inspired by the multi-label classification methods, and they pay attention to the attention sharing between labels and the personalized attention mechanism respectively, have made enough progress. *SDL* (Ben-Cohen et al. 2021) has a certain hard split on the principal vector compared with *Fast0Tag* (Zhang, Gong, and Shah 2016), which objectively improves the multi-label performance. However, *(ML)<sup>2</sup>P-Encoder* (Liu et al. 2023) starts from the perspective of the channel and uses the response for semantic encoding, achieving sub-optimal results in the two main indicators of F1-Score and mAP. Judging from the results of the GZSL task, *SDL* (Ben-Cohen et al. 2021) uses image principal vectors to generate more generalized semantics, which has absolute advantages in the GZSL environment. However, due to the design of Group Prompts, our method also learns sufficient semantic information through training. At the same time, more comprehensive global semantics also help us obtain more generalized image information, which greatly improves the ZSL and GZSL performance of the model, whether  $K = 3$  or  $K = 5$ .

Table 2 shows the performance of ours and other competitive methods on the *Open-Images-V4* test-set. For fairness, we keep the same training epoch as the comparison method. In the ZSL task, due to the limited number of ZSL labels contained in most samples, the precision indicators of all methods are low. However, our model still achieves relatively the best results, both in terms of precision and recall. In the scenario of this data set, it can be seen that the performance of *LESA* and *BiAM* has dropped significantly

Method	ZSL						GZSL							
	mAP	K=10			K=20			mAP	K=10			K=20		
		P	R	F1	P	R	F1		P	R	F1	P	R	F1
Fast0Tag (Zhang, Gong, and Shah 2016)	41.2	0.3	12.6	0.7	0.3	21.3	0.6	45.2	14.8	17.3	16.0	9.3	21.5	13.0
LESA ApC (Huynh and Elhamifar 2020)	40.7	0.6	22.9	1.2	0.4	32.4	0.9	44.9	15.7	18.3	16.9	9.6	22.3	13.5
LESA (Huynh and Elhamifar 2020)	41.7	0.7	25.6	1.4	0.5	37.4	1.0	45.4	16.2	18.9	17.4	10.2	23.9	14.3
BiAM (Narayan et al. 2021)	62.8	2.3	17.8	4.1	2.0	31.0	3.7	79.6	16.4	18.9	17.6	10.9	25.0	15.1
SDL (Ben-Cohen et al. 2021)	62.9	<b>6.1</b>	<b>47.0</b>	<b>10.7</b>	<b>4.4</b>	<b>68.1</b>	<b>8.3</b>	75.3	<b>35.3</b>	<b>40.8</b>	<b>37.8</b>	<b>23.6</b>	<b>54.5</b>	<b>32.9</b>
Gen-MLZSL (Gupta et al. 2023)	43.0	1.3	42.4	2.5	1.1	52.1	2.2	75.5	33.6	38.9	36.1	22.8	52.8	31.9
(ML) <sup>2</sup> P-Encoder (Liu et al. 2023)	<b>65.7</b>	-	-	7.5	-	-	6.5	<b>79.9</b>	-	-	27.6	-	-	24.1
Ours	<b>68.5</b>	<b>7.7</b>	<b>53.2</b>	<b>13.5</b>	<b>4.7</b>	<b>64.3</b>	<b>8.8</b>	<b>81.0</b>	<b>37.2</b>	<b>41.8</b>	<b>39.4</b>	<b>24.9</b>	<b>57.2</b>	<b>34.7</b>

Table 2: Comparison of different models for multi-label ZSL and GZSL tasks on the **Open-Images-V4** dataset. The ‘‘P’’ refers to the top-K ‘‘Precision’’ and ‘‘R’’ refers to the top-K ‘‘Recall’’. Best results are shown in **red**, and the second are shown in **blue**.

compared with *SDL*. This is mainly because as the total number of labels increases, the model’s demand for generalization performance also increases. This is what traditional multi-label classification models lack and, therefore, perform poorly. The most difficult situation occurs in GZSL, where the number of labels that need to be predicted reaches around 5,000, which is undoubtedly a severe challenge to the model. In this case, our mAP is still the best. In terms of F1-Score performance, we are still 1.6% and 1.8% ahead of *SDL*, respectively.

Model	a	b	c	d	e	f	ours
Backbone Only	✓	✓					
+GPA Module			✓	✓			✓
+GFP Module					✓	✓	✓
M=8	✓		✓		✓		✓
M=2		✓		✓		✓	
mAP ZSL	28.6	27.9	32.4	30.7	31.8	30.5	<b>33.2</b>
mAP GZSL	10.8	10.6	13.2	12.5	13.6	12.8	<b>14.7</b>

Table 3: **Ablation study** results. It shows the contribution of the different modules in *Epsilon*. All the models are performed on the test-set of the *NUS-Wide*.




Figure 3: **Hyper-parameter selection**. All the experiments are performed on the *NUS-Wide* test-set.

### Hyper-Parameters and Ablation Study

**Hyper-Parameters:** In terms of hyper-parameter selection, our model has two hyper-parameters,  $M$  and  $\lambda$ , as shown in the figure. First of all, we focus on the number of groups  $M$ . It can be seen that the model we designed does not mean that the larger the  $M$ , the better the performance. This is because once the features are grouped into too many prompts, the se-

mantic information will be extremely scattered, which will lead to the decline of mAP results. But when the number of  $M$  is too small, the semantic diversity cannot be protected. Thus, the semantic output is not comprehensive. Therefore, it can be seen from the experimental results that the model can achieve relatively best results when  $M$  is 8. For another hyper-parameter  $\lambda$ , we also conducted experimental verification. It can be seen that when  $\lambda = 0.1$ , the loss function of the model mostly includes  $\mathcal{L}_r$ , and  $\mathcal{L}_{reg}$  has the weakest influence on the loss function. When  $\lambda = 0.9$ , the loss function of the model mostly contains the regular term  $\mathcal{L}_{reg}$ . From the perspective of the mAP results caused by the change of  $\lambda$ , adding the regularization term appropriately will help the convergence of the model. From the observation in Figure 3 (b), we can discover that when  $\lambda = 0.3$ , the model can obtain relatively optimal results.

**Ablation Study:** In order to verify the effectiveness of each module of our method and to test the integrity of the method, we designed an ablation study experiment. By comparing models ‘a’ and ‘b’, ‘c’ and ‘d’, and ‘e’ and ‘f’, we can find that the impact of changing hyper-parameter  $M$  will not be affected by changes in the model structure. The experimental results of the model when  $M = 8$  are always stronger than those when  $M = 2$ . In addition, through horizontal comparison, we will find that the addition of each module has a positive effect on the performance improvement of the model’s ZSL and GZSL tasks. But they each have their areas of expertise. For example, we can compare the improvements brought by the GPA module and the GFP module to the model between ‘c’ and ‘e’. We can observe that the GPA module is obviously more beneficial to the ZSL task, while GFP is better at handling the GZSL task. This is because global features generally have stronger generalization capabilities than local group aggregation of features,

Method	P(K=3)	R(K=3)	F1(K=3)	P(K=5)	R(K=5)	F1(K=5)	mAP
<i>Logistic</i>	46.1	57.3	51.1	34.2	70.8	46.1	21.6
<i>WARP</i>	49.1	61.0	54.4	36.6	75.9	49.4	3.1
<i>WSABIE</i>	48.5	60.4	53.8	36.5	75.6	49.2	3.1
<i>Fast0Tag</i>	48.6	60.4	53.8	36.0	74.6	48.6	22.4
<i>CNN-RNN</i>	49.9	61.7	55.2	37.7	78.1	50.8	28.3
One Attention per Label	51.3	63.7	56.8	38.0	78.8	51.3	32.6
One Attention per Cluster (M = 10)	51.1	63.5	56.6	37.6	77.9	50.7	31.7
<i>LESA</i> (M = 10)	52.3	65.1	58.0	38.6	80.0	52.0	31.5
<i>BiAM</i>	-	-	59.6	-	-	53.4	47.8
<i>(ML)<sup>2</sup>P-Encoder</i>	-	-	59.8	-	-	53.8	48.0
<b>Ours</b>	<b>56.3</b>	<b>70.0</b>	<b>62.4</b>	<b>41.5</b>	<b>86.0</b>	<b>55.9</b>	<b>61.1</b>

Table 4: Models comparison on **NUS-Wide** dataset under the task of **Multi-label Classification**. Best results are shown in bold.

which is why *SDL* (Ben-Cohen et al. 2021) is slightly better than *BiAM* (Narayan et al. 2021) in GZSL performance. However, local semantics is more effective for learning seen classes, so the ZSL performance of GPA is stronger than the GFP module. From the final results, we clearly know that the best results are produced when the two modules act on the model at the same time.

### Multi-Label Classification

As a prediction model, we need to accurately classify seen classes to realize the prediction of unseen classes. Therefore, the multi-label classification ability of the model is also important. Table 4 shows the performance of our model under multi-label learning. For comparison, we added some classic multi-label classification algorithms, including *Logistic Regression* (Tsoumakas and Katakis 2007), *WSABIE* (Weston, Bengio, and Usunier 2011), *WARP* (Gong et al. 2013) and *CNN-RNN* (Wang et al. 2016), and some MLZSL models. We can see from the experimental results that our method has achieved the best results compared with the traditional multi-label learning model in both F1-Score and mAP. Compared with *BiAM* (Narayan et al. 2021) and *(ML)<sup>2</sup>P-Encoder* (Liu et al. 2023), the current best-performing methods in MLZSL, our model still has considerable advantages. This is because, compared with *BiAM*, our model acquires richer semantic information and has a stronger ability to integrate class features, which can significantly improve the model’s classification ability.

### Qualitative Performance

Our model’s qualitative performance is shown in Figure 4. For the input image, the label output by the model includes not only seen labels but also unseen labels. Scene semantics is more about the model’s perception of global information. In the first picture, humans do not realize that it is a boat sailing on the sea at first glance. However, the model first recognized ‘boat’ and ‘blue’, and then other classes related to its semantics were also easily predicted. In the second picture, judging from the scene information, our model can accurately locate the scene as an aircraft. The scene information in the third picture is more complex, but the model can still recognize ‘water’ and even abstract classes such as ‘Canada’. The above three pictures clearly show that global

semantic diversity gives the model more perspectives to interpret global information. Detailed features rely more on the integration of local semantics. The fourth image accurately identifies the character’s gender and identifies him as an actor based on his face and clothing. The success of the above identifications undoubtedly reflects that our model is very effective in both the expression of global semantics and the aggregation of local semantics.

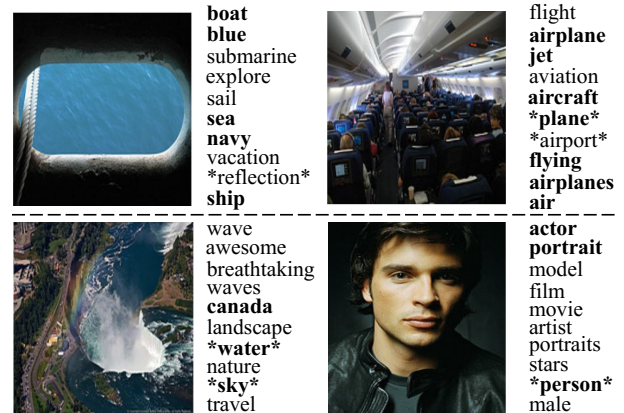


Figure 4: Top-10 labels predicted by *Epsilon* in the case of Generalized MLZSL on NUS-Wide dataset. The asterisk marks indicate unseen labels, while the bold ones indicate successfully predicted seen and unseen labels.

### Conclusion

This paper designs *Epsilon* to solve the problem of the lack of integrity of local features and rough processing of global features, resulting in the loss of visual-semantic comprehensiveness. In response to the above problems, we first designed the GPA module to improve the integrity of local semantic information by grouping and aggregating spatial features. Then, the GFP module will be used to obtain more diverse and comprehensive global features. Finally, the global and local features are combined as complete semantic information of an image for training. On two publicly available large MLZSL datasets *NUS-Wide* and *Open-Images-V4*, our experimental results far outperform other state-of-the-art models.

## Acknowledgments

This research was supported by funding from the Hong Kong RGC General Research Fund (No. 15221123 and 15216424), PolyU Internal Fund (No. P0043932)

## References

- Akata, Z.; Malinowski, M.; Fritz, M.; and Schiele, B. 2016. Multi-cue zero-shot learning with strong supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 59–68.
- Ben-Cohen, A.; Zamir, N.; Ben-Baruch, E.; Friedman, I.; and Zelnik-Manor, L. 2021. Semantic diversity learning for zero-shot multi-label classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 640–650.
- Chen, Z.-M.; Wei, X.-S.; Wang, P.; and Guo, Y. 2019. Multi-label image recognition with graph convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5177–5186.
- Cheng, X.; Lin, H.; Wu, X.; Yang, F.; Shen, D.; Wang, Z.; Shi, N.; and Liu, H. 2021. MITr: Multi-label Classification with Transformer. *arXiv preprint arXiv:2106.06195*.
- Chua, T.-S.; Tang, J.; Hong, R.; Li, H.; Luo, Z.; and Zheng, Y. 2009. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*, 1–9.
- Deutsch, S.; Kolouri, S.; Kim, K.; Owechko, Y.; and Soatto, S. 2017. Zero shot learning via multi-scale manifold regularization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7112–7119.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Durand, T.; Mehra, N.; and Mori, G. 2019. Learning a deep convnet for multi-label classification with partial labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 647–657.
- Feng, L.; An, B.; and He, S. 2019. Collaboration based multi-label learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3550–3557.
- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Ranzato, M.; and Mikolov, T. 2013. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26.
- Gao, B.-B.; and Zhou, H.-Y. 2021. Learning to discover multi-class attentional regions for multi-label image recognition. *IEEE Transactions on Image Processing*, 30: 5920–5932.
- Gong, Y.; Jia, Y.; Leung, T.; Toshev, A.; and Ioffe, S. 2013. Deep convolutional ranking for multilabel image annotation. *arXiv preprint arXiv:1312.4894*.
- Guo, J.; and Guo, S. 2020. A novel perspective to zero-shot learning: Towards an alignment of manifold structures via semantic feature expansion. *IEEE Transactions on Multimedia*, 23: 524–537.
- Gupta, A.; Narayan, S.; Khan, S.; Khan, F. S.; Shao, L.; and Van De Weijer, J. 2023. Generative multi-label zero-shot learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Huynh, D.; and Elhamifar, E. 2020. A shared multi-attention framework for multi-label zero-shot learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8776–8786.
- Kim, J.-H.; Jun, J.; and Zhang, B.-T. 2018. Bilinear attention networks. *arXiv preprint arXiv:1805.07932*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kodirov, E.; Xiang, T.; and Gong, S. 2017. Semantic auto-encoder for zero-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3174–3183.
- Lee, C.-W.; Fang, W.; Yeh, C.-K.; and Wang, Y.-C. F. 2018. Multi-label zero-shot learning with structured knowledge graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1576–1585.
- Liu, Z.; Guo, S.; Lu, X.; Guo, J.; Zhang, J.; Zeng, Y.; and Huo, F. 2023.  $\mathcal{L}_2$  P-Encoder: On Exploration of Channel-Class Correlation for Multi-Label Zero-Shot Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23859–23868.
- Mensink, T.; Gavves, E.; and Snoek, C. G. 2014. Costa: Co-occurrence statistics for zero-shot classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2441–2448.
- Narayan, S.; Gupta, A.; Khan, S.; Khan, F. S.; Shao, L.; and Shah, M. 2021. Discriminative region-based multi-label zero-shot learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8731–8740.
- Norouzi, M.; Mikolov, T.; Bengio, S.; Singer, Y.; Shlens, J.; Frome, A.; Corrado, G. S.; and Dean, J. 2014. Zero-shot learning by convex combination of semantic embeddings. In *2nd International Conference on Learning Representations, ICLR 2014*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; Rabinovich, A.; et al. 2015. Going deeper with convolutions. *Cvpr*.
- Tsoumakas, G.; and Katakis, I. 2007. Multi-label classification: An overview. *International Journal of Data Warehousing and Mining (IJDWM)*, 3(3): 1–13.
- Veit, A.; Alldrin, N.; Chechik, G.; Krasin, I.; Gupta, A.; and Belongie, S. 2017. Learning from noisy large-scale datasets with minimal supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 839–847.

Wang, J.; Yang, Y.; Mao, J.; Huang, Z.; Huang, C.; and Xu, W. 2016. Cnn-rnn: A unified framework for multi-label image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2285–2294.

Weston, J.; Bengio, S.; and Usunier, N. 2011. Wsabie: Scaling up to large vocabulary image annotation. In *Twenty-Second International Joint Conference on Artificial Intelligence*.

Xian, Y.; Schiele, B.; and Akata, Z. 2017. Zero-shot learning—the good, the bad and the ugly. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4582–4591.

Yu, H.-F.; Jain, P.; Kar, P.; and Dhillon, I. 2014. Large-scale multi-label learning with missing labels. In *International conference on machine learning*, 593–601. PMLR.

Zhang, Y.; Gong, B.; and Shah, M. 2016. Fast zero-shot image tagging. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5985–5994. IEEE.