

# Cross-Modality Attention with Semantic Graph Embedding for Multi-Label Classification

Renchun You,<sup>1\*</sup> Zhiyao Guo,<sup>2\*</sup> Lei Cui,<sup>3\*</sup> Xiang Long,<sup>1</sup> Yingze Bao,<sup>1</sup> Shilei Wen<sup>1†</sup>

<sup>1</sup>Baidu VIS

<sup>2</sup>Computer Science Department, Xiamen University, China

<sup>3</sup>Department of Computer Science and Technology, Tsinghua University, China

{yourenchun, longxiang, wenshilei}@baidu.com, {guozhiyao45, baoyingze}@gmail.com, cuil19@mails.tsinghua.edu.cn

## Abstract

Multi-label image and video classification are fundamental yet challenging tasks in computer vision. The main challenges lie in capturing spatial or temporal dependencies between labels and discovering the locations of discriminative features for each class. In order to overcome these challenges, we propose to use cross-modality attention with semantic graph embedding for multi-label classification. Based on the constructed label graph, we propose an adjacency-based similarity graph embedding method to learn semantic label embeddings, which explicitly exploit label relationships. Then our novel cross-modality attention maps are generated with the guidance of learned label embeddings. Experiments on two multi-label image classification datasets (MS-COCO and NUS-WIDE) show our method outperforms other existing state-of-the-arts. In addition, we validate our method on a large multi-label video classification dataset (YouTube-8M Segments) and the evaluation results demonstrate the generalization capability of our method.

## 1 Introduction

Multi-label image classification (MLIC) and multi-label video classification (MLVC) are important tasks in computer vision, where the goal is to predict a set of categories present in an image or a video. Compared with single-label classification (e.g. assigns one label to an image or video), multi-label classification is more useful in many applications such as internet search, security surveillance, robotics, etc. Since MLIC and MLVC are very similar tasks, in the following technical discussion we will mainly focus on MLIC, whose conclusions can be migrated to MLVC naturally.

Recently, single-label image classification has achieved great success thanks to the evolution of deep Convolutional Neural Networks (CNN) (He et al. 2016; Huang et al. 2017; Simonyan and Zisserman 2014; Szegedy et al. 2016). Single-label image classification can be naively extended to MLIC tasks by treating the problem as a series of single-label classification tasks. However, such naive extension

usually provides poor performance, since the semantic dependencies among multiple labels are ignored, which are especially important for multi-label classification. Therefore, a number of prior works aim to capture the label relations by Recurrent Neural Networks (RNN). However, these methods do not model the explicit relationships between semantic labels and image regions, thus they lack the capacity of sufficient exploitation of the spatial dependency in images.

An alternative solution for MLIC is to introduce object detection techniques. Some methods (Wei et al. 2014; Zhang et al. 2018; Hao et al. 2016) extract region proposals using extra bounding box annotations, which are much more expensive to label than simple image level annotations. Many other methods (Wang et al. 2017; Zhu et al. 2017) apply attention mechanism to automatically focus on the regions of interest. However, the attentional regions are learned only with image-level supervision, which lacks explicit semantic guidance.

To address above issues, we argue that an effective model for multi-label classification should reach two capacities: (1) capturing semantic dependencies among multiple labels in terms of spatial context; (2) locating regions of interest with more semantic guidance.

In this paper, we propose a novel cross-modality attention network associated with graph embedding, so as to simultaneously search for discriminative regions and label spatial semantic dependencies. Firstly, we introduce a novel Adjacency-based Similarity Graph Embedding (ASGE) method which captures the rich semantic relations between labels. Secondly, the learned label embedding will guide the generation of attentional regions in terms of cross-modality guidance, which is referred to as Cross-modality Attention (CMA) in this paper. Compared with traditional self-attention methods, our attention explicitly introduces the rich label semantic relations. Benefiting from the CMA mechanism, our attentional regions are more meaningful and discriminative. Therefore they capture more useful information while suppressing the noise or background information for classification. Furthermore, the spatial context dependencies of labels will be captured, which further improve the performance in MLIC.

The major contributions of this paper are briefly summa-

\*These authors share first authorship.

†Corresponding author.

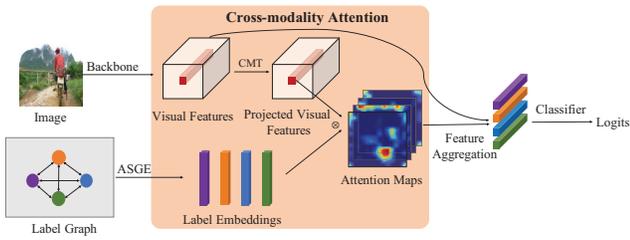


Figure 1: The overall framework of our model for MLIC task. The label embeddings are obtained by ASGE module. The visual features are first extracted by backbone network and then projected to semantic space to get projected visual features through CMT module. The learned label embeddings and projected visual features are together fed into CMA module to generate the category-wise attention maps, each of which is utilized to weightedly average the visual features and generate category-wise aggregated feature. Finally, the classifier is applied for final prediction.

alized as follows:

- We propose an ASGE method to learn semantic label embedding and exploit label correlations explicitly.
- We propose a novel attention paradigm, namely cross-modality attention, where the attention maps are generated by leveraging more prior semantic information, resulting in more meaningful attention maps.
- A general framework combining CMA and ASGE module, as shown in Fig.1 and Fig.2, is proposed for multi-label classification, which can capture dependencies between spatial and semantic space and discover the location of discriminative feature effectively. We evaluate our framework on MS-COCO dataset and NUS-WIDE dataset for MLIC task, and new state-of-the-art performances are achieved on both of them. We also evaluate our proposed method on YouTube-8M dataset for MLVC, which also achieves remarkable performances.

## 2 Related Works

The task of MLIC has attracted an increasing interest recently. The easiest way to address this problem is to treat each category independently, then the task can be directly converted into a series of binary classification tasks (Chen et al. 2019). However, such techniques are limited by without considering the relationships between labels.

Several approaches have been applied to model the correlations between labels. Read et al. (2011) extends the multi-label classification by training the chain binary-classifiers and introducing the correlations between labels by inputting the previously predicted labels. Some other works (Li, Zhao, and Guo 2014; Wang et al. 2016; Chen et al. 2018; Li et al. 2016) formulate the task as a structural inference problem based on probabilistic graphical models. Besides, the latest work (Chen et al. 2019) explores the label dependencies by graph convolutional network. However, none of aforementioned methods consider the associations between

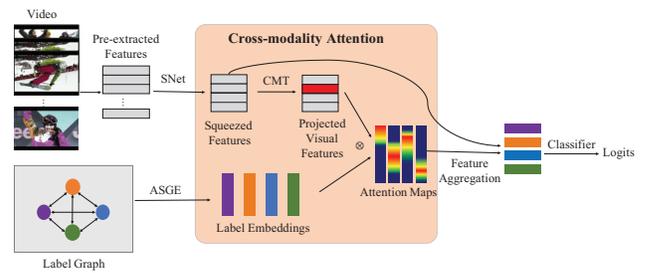


Figure 2: The overall framework of the model for MLVC task. The input is pre-extracted features instead of raw video, and the pre-extracted features are processed through SNet to extract visual features. Other parts are quite similar to that of MLIC.

semantic labels and image contents, and the spatial contexts of images have not been sufficiently exploited.

In MLIC task, visual concepts are highly related with local image regions. To explore information in local regions better, some works (Wei et al. 2014; Hao et al. 2016) introduce region proposal techniques to focus on informative regions. Wei et al. (2014) extracts an arbitrary number of object hypotheses, then inputs them into the shared CNN and aggregates the output with max pooling to obtain the ultimate multi-label predictions. Hao et al. (2016) introduces local information provided by generated proposals to boost the discriminative power of feature extraction. Although above methods have used region proposals to enhance feature representation, they are still limited by requiring extra object-level annotations and without considering the dependencies between objects.

Alternatively, Wang et al. (2017) discovers the attentional regions corresponding to multiple semantic labels by spatial transformer network and captures the spatial dependencies of the regions by Long Short-Term Memory (LSTM). Analogously, Zhu et al. (2017) proposes the spatial regularization network to generate label-related attention maps and capture the latent relationships by attention maps implicitly. The advantage of above attention approaches is that no additional step of obtaining region proposal is needed. Nevertheless, the attentional regions are learned only with image-level supervision, which lacks of explicit semantic guidance. While in this paper, the semantic guidance is introduced to the generation of attention maps by leveraging label semantic embeddings, which improves the prediction performance significantly.

In this paper, the label semantic embeddings are learned by graph embedding, which is a technique aiming to learn representation of graph-structured data. The approaches of graph embedding mainly contain matrix factorization-based (Cao, Lu, and Xu 2015), random walk-based (Perozzi, Al-Rfou, and Skiena 2014) and neural network-based methods (Wang, Cui, and Zhu 2016; Zhu et al. 2018). A main assumption of these approaches is the embeddings of adjacent nodes on the graph are similar, while in our task, we also require embeddings of non-adjacent nodes are mutually ex-

clusive from each other. Therefore, we propose an ASGE method, which can further separate the embeddings of non-adjacent nodes.

MLVC is similar to MLIC, but it involves additional temporal relationships (Gan et al. 2015; Long et al. 2018a; Campos et al. 2017; Arandjelovic et al. 2016; Wu, Ma, and Hu 2017; Long et al. 2018b). It has been applied to many applications such as emotion recognition (Kahou et al. 2016), human activity understanding (Caba Heilbron et al. 2015), and event detection (Xu, Yang, and Hauptmann 2015). In this paper, we validate our proposed method both in MLIC and MLVC task, and achieve remarkable performance.

### 3 Approach

The overall frameworks of our approach for MLIC and MLVC are shown in Fig.1 and Fig.2 respectively. The pipeline includes several stages: Firstly, the label graph is taken as the input of ASGE module to learn label embeddings which encode the semantic relationships between labels. Secondly, the learned label embeddings and visual features will be fed together into the CMA module to obtain category-wise attention maps. Finally, the category-wise attention maps are used to weightedly average the visual features for each category. We will describe our two key components ASGE and CMA in detail.

#### 3.1 Adjacency-based Similarity GE

The relationships between labels play a crucial role in multi-label classification task as discussed in section 1. However, how to express such relationships is an open issue to be solved. Our intuition is that the co-occurrence properties between labels can be described as joint probability, which is suitable for modeling the label relationships. Nevertheless, the joint probability is easy to suffer from the influence of class imbalance. Instead, we utilize the conditional probability between labels to solve this issue, which is obtained by normalizing the joint probability through dividing by marginal probability. Based on this, it is possible to construct a label graph where the labels are nodes and the conditional probability between the labels is edge weight. Inspired by the popular applications of graph embedding method in natural language processing (NLP) tasks, where the learned label embeddings are entered into the network as additional information, we propose a novel ASGE method to encode the label relationships.

We formally define the graph as  $\mathcal{G} = (\mathbf{V}, v\mathbf{C})$ , where  $v\mathbf{V} = \{v_1, v_2, \dots, v_N\}$  represents the set of  $N$  nodes and  $\mathbf{C}$  represents the edges. The adjacency matrix  $\mathbf{A} = \{A_{ij}\}_{i,j=1}^N$  of graph  $\mathcal{G}$  contains non-negative weights associated with each edge. Specifically,  $\mathbf{V}$  is the set of labels and  $\mathbf{C}$  is the set of connections between any two labels, and the adjacency matrix  $\mathbf{A}$  is the conditional probability matrix by setting  $A_{ij} = P(v_i|v_j)$ , where  $P$  is calculated through training set. Since  $P(v_i|v_j) \neq P(v_j|v_i)$ , namely  $A_{ij} \neq A_{ji}$ , in order to facilitate a better optimization, we symmetrize  $\mathbf{A}$  by

$$\mathbf{A}' = \frac{1}{2} (\mathbf{A} + \mathbf{A}^\top). \quad (1)$$

To capture the label correlations defined by the graph structure, we apply a neural network to map the one-hot embedding of each label  $\mathbf{o}_i$  to semantic embedding space and produce the label embedding

$$\mathbf{e}_i = \Phi(\mathbf{o}_i), \quad (2)$$

where  $\Phi$  denotes the neural network which consists of three fully-connected layers followed by Batch Normalization (BN) and ReLU activation. Our goal is to achieve the optimal label embedding set  $\mathbf{E} = \{\mathbf{e}_i\}_{i=0}^N$ , where  $\mathbf{e}_i \in \mathbb{R}^{C_e}$ . Such that  $\cos(\mathbf{e}_i, \mathbf{e}_j)$  is close to  $A'_{ij}$  for all  $i, j$ , where  $\cos(\mathbf{e}_i, \mathbf{e}_j)$  denotes the cosine similarity between  $\mathbf{e}_i$  and  $\mathbf{e}_j$ . Thereby, the objective function is defined as follows:

$$\mathcal{L}_{ge} = \sum_{i=1}^N \sum_{j=1}^N \left( \frac{\mathbf{e}_i^\top \mathbf{e}_j}{\|\mathbf{e}_i\| \|\mathbf{e}_j\|} - A'_{ij} \right)^2, \quad (3)$$

where  $\mathcal{L}_{ge}$  denotes the loss of our graph embedding.

**Optimization Relaxation.** In order to optimize Eq.3, the cosine similarity  $\cos(\mathbf{e}_i, \mathbf{e}_j)$  are required to be close to the corresponding edge weight  $A'_{ij}$  for all  $i, j$ . However, it is hard to satisfy this strict constraint, especially when the graph is large and sparse. In order to address this problem, a hyperparameter  $\alpha$  is introduced to the Eq.3 to relax the optimization. The new objective function is as follows:

$$\mathcal{L}_{ge} = \sum_{i=1}^N \sum_{j=1}^N \sigma_{ij} \cdot \left( \frac{\mathbf{e}_i^\top \mathbf{e}_j}{\|\mathbf{e}_i\| \|\mathbf{e}_j\|} - A'_{ij} \right)^2, \quad (4)$$

where  $\sigma_{ij}$  is an indicator function:

$$\sigma_{ij} = \begin{cases} 0, & A_{ij} < \alpha \text{ and } \frac{\mathbf{e}_i^\top \mathbf{e}_j}{\|\mathbf{e}_i\| \|\mathbf{e}_j\|} < \alpha \\ 1, & \text{otherwise} \end{cases}. \quad (5)$$

By adding this relaxation, it only needs to make the embedding pairs  $(\mathbf{e}_i, \mathbf{e}_j)$  be away instead of strictly enforcing  $\cos(\mathbf{e}_i, \mathbf{e}_j)$  to be  $A'_{ij}$  when  $A_{ij} < \alpha$ , thus focusing more on the strong relationships between labels and reducing the difficulty of the optimization.

#### 3.2 CMA for Multi-label Classification

We formally define the multi-label classification task as a mapping function  $F : \mathbf{x} \rightarrow \mathbf{y}$ , where  $\mathbf{x}$  denotes a input image or video,  $\mathbf{y} = [y_1, y_2, \dots, y_N]$  denotes corresponding labels,  $N$  is the total number of categories and  $y_n \in \{0, 1\}$  denotes whether the label is assigned to the image or video.

For multi-label classification, we propose an novel attention mechanism, named cross-modality attention, which uses semantic embeddings to guide spatial or temporal integration of visual features. The semantic embeddings here are the label embedding set  $\mathbf{E} = \{\mathbf{e}_i\}_{i=0}^N$  achieved by ASGE and the visual features  $\mathbf{I} = \psi(\mathbf{x})$  are extracted by backbone neural networks  $\psi$ . Note that for different tasks, we only need to apply different backbones to extract visual features, and the rest part of the framework is completely generic for both tasks.

**Backbone.** In the MLIC task, we apply ResNet-101 network to extract the last convolutional feature map as the visual features. Additionally, we use an  $1 \times 1$  convolution to reduce the dimension, and obtain final visual feature map  $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$ , where  $H \times W$  is the spatial resolution of the last feature map and  $C$  is the number of channels.

In the MLVC task, frame level features  $\mathbf{x}$  are pre-extracted by an Inception network and then processed by PCA with whitening. Considering the meaningful and discriminative information of video is derived from some pivotal frames while others may be redundant, we apply a Squeezing Network (SNet) to squeeze the temporal dimensions. The SNet is built on 4 successive 1D convolution and pooling layers. We can obtain the final visual feature  $\mathbf{I} = f_{SNet}(\mathbf{x})$ , where  $\mathbf{I} \in \mathbb{R}^{T \times C}$  and  $T$  is the temporal resolution of the final feature map and  $C$  is the number of channels.

**Cross-Modality Attention.** The learned label embeddings by ASGE compose a semantic embedding space, while the extracted features from CNN Backbone define a visual feature space. Our goal is to let semantic embeddings guide the generation of attention maps. However, semantic embedding space and visual feature space exist a semantic gap because of modality difference. In order to measure the compatibility between different modalities, we first learn a mapping function from the visual feature space to the semantic embedding space, then the compatibility can be measured by a cosine similarity between projected visual feature and semantic embedding, namely cross-modality attention. Formal definition is introduced as follows.

Firstly, we project the visual feature to semantic space by a Cross-Modality Transformer (CMT) module, which is built with several  $1 \times 1$  convolution layers followed by a BN and a ReLU activation.

$$\mathbf{I}_s = f_{cmt}(\mathbf{I}) \quad (6)$$

where  $\mathbf{I}_s \in \mathbb{R}^{M \times C_e}$  ( $M = W \times H$  for image and  $M = T$  for video),  $f_{cmt}$  denotes the map function of the CMT module. The category-specific cross-modality attention map  $z_k^i$  is yielded by calculating the cosine similarity between label embedding  $\mathbf{e}_k$  and projected visual feature vector  $\mathbf{I}_s^i$  at location  $i$  of  $\mathbf{I}_s$ :

$$z_k^i = \text{ReLU} \left( \frac{\mathbf{I}_s^i \top \mathbf{e}_k}{\|\mathbf{I}_s^i\| \|\mathbf{e}_k\|} \right). \quad (7)$$

The category-specific attention map  $z_k^i$  is then normalized to:

$$a_k^i = \frac{z_k^i}{\sum_{i=1}^M z_k^i}. \quad (8)$$

For each location  $i$ , if the CMA mechanism generates a high positive value, it can be interpreted as the location  $i$  is highly semantic related to label embedding  $k$  or relative more important than other locations, thus the model needs to focus on location  $i$  when considering category  $k$ . Then the category-specific cross-modality attention map is used to weightedly average the visual feature vectors for each category:

$$\mathbf{h}_k = \sum_{i=1}^M \alpha_k^i \mathbf{I}^i, \quad (9)$$

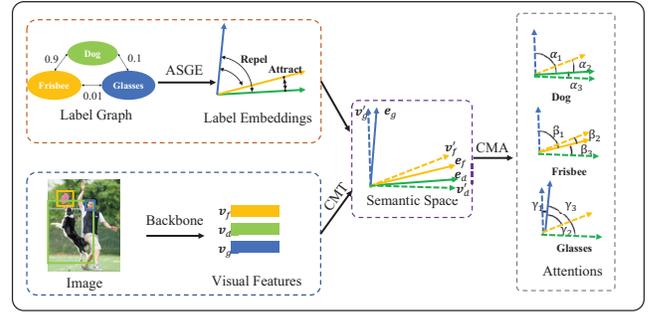


Figure 3: Illustration of latent spatial dependency. The different colors indicate different categories. The solid arrows represent the learned label embeddings, expressed as  $\mathbf{e}$ , while the dotted arrows represent the projected visual features through CMT module, expressed as  $\mathbf{v}'$ . The angles between label embeddings and projected visual features, namely  $\alpha$ ,  $\beta$  and  $\gamma$ , represent the category-wise attention scores. For a detailed discussion, refer to section 3.2.

where  $\mathbf{h}_k$  is the final feature vector for label  $k$ . Then  $\mathbf{h}_k$  is fed into the fully-connected layers for estimating probability of category  $k$ :

$$y_k^* = \sigma(\mathbf{w}_k \top \mathbf{h}_k + b), \quad (10)$$

where  $\mathbf{w}_k \in \mathbb{R}^C$  and  $b$  are learnable parameters.  $y_k^*$  is the predicted probability for label  $k$ . For convenience, we denote the calculation of whole CMA module as  $y_k^* = f_{cma}(\mathbf{I}, \mathbf{E})$ .

Compared with general single attention map method, where the attention map is shared by all categories, our CMA module benefits in two ways: Firstly, our category-wise attention map is related to image regions corresponding to category  $k$ , thus better learn category-related regions. Secondly, with the guidance of label semantic embeddings, the discovered attentional regions can be better match with the annotated semantic labels.

The potential advantage of our framework is to capture the latent spatial dependency, which is helpful for visual ambiguous labels. As shown in Fig.3, we consider *frisbee* as an example to explain the spatial dependency. Firstly, The ASGE module learns label embeddings through the label graph, which encodes the label relationships. Since the *dog* and *frisbee* are often co-exist while *glasses* not, therefore the label embeddings of *dog* and *frisbee* are close to each other and far away from *glasses*'s, namely  $\mathbf{e}_d \approx \mathbf{e}_f \neq \mathbf{e}_g$ . The optimization procedure during training will enforce the cosine similarity between visual feature and the corresponding label embedding become high, in other words, the  $\cos(\mathbf{e}_d, \mathbf{v}'_d)$ ,  $\cos(\mathbf{e}_g, \mathbf{v}'_g)$  and  $\cos(\mathbf{e}_f, \mathbf{v}'_f)$  will be large. Since  $\mathbf{e}_d \approx \mathbf{e}_f \neq \mathbf{e}_g$ ,  $\cos(\mathbf{e}_f, \mathbf{v}'_d)$  will also be large, while  $\cos(\mathbf{e}_f, \mathbf{v}'_g)$  will be small. And the final feature representation of *frisbee* is  $\mathbf{h}_f = \beta_1 \mathbf{v}_g + \beta_2 \mathbf{v}_f + \beta_3 \mathbf{v}_d$ , where  $\beta_1 = \cos(\mathbf{e}_f, \mathbf{v}'_g)$ ,  $\beta_2 = \cos(\mathbf{e}_f, \mathbf{v}'_f)$ ,  $\beta_3 = \cos(\mathbf{e}_f, \mathbf{v}'_d)$ . Thus, the recognition of *frisbee* is depending on the semantic related label *dog* and not related to label *glasses*, indicating that our model is capable of capturing spatial dependencies. Specially, considering that the *frisbee* is a hard case to be

recognized,  $\beta_2$  will be small. Fortunately, the  $\beta_3$  may still be large, so the visual information of *dog* will be a helpful context to aid in the recognition of label *frisbee*.

**Multi-Scale CMA.** Single-scale feature representation may not be sufficient for multiple objects from different scales. It is noteworthy that the calculation of attention involves the label embedding slides densely over all locations of the feature map, in other words, the spatial resolution of feature map may effect on attention result. Our intuition is that the low-resolution feature maps have more representational capacity for small objects while high-resolution is opposite. The design of CMA mechanism makes it can be naturally applied to multi-scale feature maps via a score fusion strategy. Specially, we extract a set of feature maps  $\{I_1, I_2, \dots, I_L\}$  and the final predicted probability of multi-scale CMA is

$$y_k^* = \frac{1}{L} \sum_{l=1}^L f_{cma}(I_l, E). \quad (11)$$

**Training Loss.** Finally we define our object function for multi-label classification as follows

$$L(\theta) = - \sum_{k=1}^N w_k [y_k \cdot \log(y_k^*) + (1 - y_k) \cdot (1 - \log(y_k^*))] \\ w_k = y_k \cdot e^{\beta(1-p_k)} + (1 - y_k) \cdot e^{\beta p_k}, \quad (12)$$

Where  $w_k$  is used to alleviate the class imbalance,  $\beta$  is a hyperparameter and  $p_k$  is the ratio of label  $k$  in the training set.

## 4 Experiments

To assess our model, we perform experiments on two benchmark multi-label image recognition datasets (MS-COCO (Lin et al. 2014) and NUS-WIDE (Chua et al. 2009)). We also validate the effectiveness of our model on one multi-label video recognition dataset (YouTube-8M Segments), and the results demonstrate the extensibility of our method. In this section, we will introduce the results on MLIC and MLVC respectively.

### 4.1 Multi-label Image Classification

**Implementation Details.** In ASGE module, the dimensions of the three hidden layers and label embeddings are all set as 256. The optimization relaxation is not applied here since the label graph is relatively small. The optimizer is Stochastic Gradient Descent (SGD) with momentum 0.9 and the initial learning rate is 0.01. In the classification part, the input image is randomly cropped and resized to  $448 \times 448$  with random horizontal flip for augmentation. The batch size is set as 64. The optimizer is SGD with momentum 0.9. Weight decay is  $10^{-5}$ . The initial learning rate is 0.01 and decays by a factor 10 every 30 epochs. And the hyperparameter  $\beta$  in the Eq.12 is 0. in MS-COCO dataset and 0.4 in NUS-WIDE dataset. Based on this setup, we implement two models: CMA and Multi-Scale CMA(MS-CMA). The MS-CMA model uses three scale features  $I_1 \in \mathbb{R}^{28 \times 28 \times 1024}$ ,  $I_2 \in \mathbb{R}^{14 \times 14 \times 2048}$  from

ResNet-101 backbone and  $I_3 \in \mathbb{R}^{7 \times 7 \times 512}$  obtained by applying a residual block on  $I_2$ . While the CMA model only uses  $I_2$ .

**Evaluation Metrics.** We use the same evaluation metrics as other works (Wang et al. 2017), which are the per-category and overall metrics: precision (CP and OP), recall (CR and OR) and F1 (CF1 and F1). In addition, we also calculate the mean average precision (mAP), which is relatively more important than other metrics, and we mainly focus on the performance of mAP.

**Results on MS-COCO Dataset.** The MS-COCO dataset is widely used in MLIC task. It contains 122,218 images with 80 labels and almost 2.9 labels per image. We divide the dataset into two parts: 82,081 images for training and 40,137 images for testing, according to the officially provided division criteria.

We compare with the currently published state-of-the-art methods, including CNN-RNN(Wang et al. 2016), RNN-Attention(Wang et al. 2017), Order-Free RNN(Chen et al. 2018), ML-ZSL (Lee et al. 2018), SRN (Zhu et al. 2017) and Multi-Evidence (Ge, Yang, and Yu 2018). Besides, we run the source code released by ML-GCN(Chen et al. 2019) to train and get the results for comparison. The quantitative results of CMA and MS-CMA model are shown in Table 1. Our two models both perform better than the state-of-the-art methods over almost all metrics. Specially, our MS-CMA model achieves better performance than the CMA model, demonstrating the multi-scale attentions yield performance improvement.

**Results on NUS-WIDE Dataset.** The NUS-WIDE is a web dataset including 269,648 images and 5018 labels from the Flickr. After removing the noise and the rare labels, there are 1000 categories left. The images are further manually annotated into 81 concepts with 2.4 concepts per image on average. We follow the split used in (Liu et al. 2018), i.e. 150,000 images for training and 59,347 for testing after removing the images without any labels.

In this dataset, we compare with the current state-of-the-art models, including CNN-RNN (Wang et al. 2016), CNN-SREL-RNN (Liu et al. 2017), CNN-LSEP (Li, Song, and Luo 2017), Order-Free RNN(Chen et al. 2018), ML-ZSL (Lee et al. 2018), S-CLs(Liu et al. 2018), Attention transfer(Zagoruyko and Komodakis 2016) and FitsNet(Romero et al. 2014).

The quantitative results are shown in the Table 2. The comparison results are similar to MS-COCO's. Our CMA and MS-CMA perform better than state-of-the-art methods on most metrics. The mAP metric of our MS-CMA, which is mostly concerned, exceeds the previous state-of-the-art result by 1.3%. We observe that the average edge weight per label of NUS-WIDE is 3.3, while that of MS-COCO is 3.9. This shows that the label graph of MS-COCO is denser than NUS-WIDE. And the performance gain of NUS-WIDE is slightly less obvious than that of MS-COCO. These observations indicate that richer label relationships may bring performance improvement.

Methods	All							Top-3					
	mAP	CP	CR	CF1	OP	OR	OF1	CP	CR	CF1	OP	OR	OF1
CNN-RNN (2016)	61.2	-	-	-	-	-	-	66.0	55.6	60.4	69.2	66.4	67.8
CNN-LSEP (2017)	73.5	56.4	62.9	76.3	61.8	68.3	-	-	-	-	-	-	-
CNN-SREL-RNN (2017)	67.4	59.8	63.4	76.6	68.7	72.4	-	-	-	-	-	-	-
RNN-Attention (2017)	-	-	-	-	-	-	-	79.1	58.7	67.4	82.0	63.0	72.0
Order-Free RNN (2018)	-	-	-	-	-	-	-	71.6	54.8	62.1	74.2	62.2	67.7
ML-ZSL (2018)	-	-	-	-	-	-	-	74.1	64.5	69.0	-	-	-
SRN (2017)	77.1	81.6	65.4	71.2	82.7	69.9	75.8	85.2	58.8	67.4	87.4	62.5	72.9
S-CLs (2018)	74.6	-	-	69.2	-	-	74.0	-	-	66.8	-	-	72.7
Multi-Evidence (2018)	-	80.4	70.2	74.9	<u>85.2</u>	72.5	78.4	84.5	62.2	70.6	89.1	64.3	74.7
ML-GCN (2019)	82.4	82.1	73.1	77.3	83.7	76.3	79.9	<u>87.2</u>	<u>64.6</u>	<u>74.2</u>	89.1	66.7	76.3
<b>CMA</b>	<u>82.8</u>	<b>83.2</b>	72.5	<u>77.5</u>	<b>86.5</b>	75.9	80.9	86.6	64.4	73.8	<b>90.8</b>	66.9	77.0
<b>MS-CMA</b>	<b>83.8</b>	<u>82.9</u>	<b>74.4</b>	<b>78.4</b>	84.4	<b>77.9</b>	<b>81.0</b>	<b>88.2</b>	<b>65.0</b>	<b>74.9</b>	<u>90.2</u>	<b>67.4</b>	<b>77.1</b>

Table 1: Comparisons with state-of-the-art methods on the MS-COCO dataset. We report two our proposed model: **CMA** and **MS-CMA**. The bold numbers indicate the best results in different metrics, while the underlined numbers indicate the suboptimal results.

Methods	All			Top-3	
	mAP	CF1	OF1	CF1	OF1
CNN-RNN (2016)	56.1	-	-	34.7	55.2
CNN-LSEP (2017)	-	52.9	70.8	-	-
CNN-SREL-RNN (2017)	-	52.7	70.9	-	-
Order-Free RNN (2018)	-	-	-	54.7	<u>70.2</u>
ML-ZSL (2018)	-	-	-	45.7	-
Attention transfer (2016)	57.6	55.2	70.3	51.7	68.8
FitsNet (2014)	57.4	54.9	70.4	51.4	68.6
S-CLs (2018)	60.1	58.7	<u>73.7</u>	53.8	<b>71.1</b>
<b>CMA</b>	<u>60.8</u>	<u>60.4</u>	<u>73.7</u>	55.5	70.0
<b>MS-CMA</b>	<b>61.4</b>	<b>60.5</b>	<b>73.8</b>	<b>55.7</b>	69.5

Table 2: Comparisons with state-of-the-art methods on the NUS-WIDE dataset.

**Ablation Study.** In this section, we expect to answer the following questions:

- Comparing with the backbone (ResNet-101) model, does our CMA model improve significantly?
- Does our proposed CMA mechanism have an advantage over the general self-attention methods?
- Can the CMA extend to multi-scale and bring performance improvement?
- Is our ASGE more advantageous than other embedding methods, e.g. Word2vec?

To answer these questions, we conduct some ablation studies on the MS-COCO dataset, as shown in table 3. Firstly, we investigate how CMA contributes to mAP. It is obvious to see that the vanilla ResNet-101 achieves 79.9% mAP, while increases to 82.8% when CMA module is added. This result shows the significant effectiveness of the CMA mechanism. Secondly, we implement a general self-attention method by replacing Eq.7 with  $z^i = \sigma(f_{conv}(I^i))$ , where  $f_{conv}$  denotes the map function of  $1 \times 1$  convolution layer. Our CMA mechanism performs better than general self-attention mechanism by achieving 1.7% mAP improvement, which indicates that the label semantic embeddings

Methods	mAP
ResNet-101 (2016)	79.9
Self-attention	81.1
W2V-MS-CMA	82.5
<b>CMA</b>	82.8
<b>MS-CMA</b>	83.8

Table 3: Comparison of mAP with several models on the MS-COCO dataset.

guided attention mechanism is superior to the general self-attention due to introducing much more prior information. Thirdly, expanding our CMA mechanism to multiple scales could obtain about 1% improvement. This result demonstrates that our attention mechanism is well adapted for multi-scale features. Finally, we compared our ASGE with other embedding methods, in this paper, we take Word2vec as an example, which is a group of related models used to produce word embeddings. Specially, we view the label set in each image as a single sentence, and the window size in Word2vec is set as the length of longest sentence to eliminate the influence of label order. The experiment results show our ASGE based MS-CMA performs better than Word2vec based MS-CMA (represented as W2V-MS-CMA) by 1.3% improvement. In our ASGE, label relationships are explicitly represented by adjacency matrix which is treated as a direct optimization target. Instead, Word2vec implicitly encodes label relationships in a data-driven manner without directly optimizing label relationships. Therefore, our ASGE will capture label relationships much better.

**Visualization and Analysis.** In this section, we visualize the learned attention maps to illustrate the ability of exploiting discriminative or meaningful regions and capturing the spatial semantic dependencies.

We show the attention visualization examples in Fig.4. The three rows show the category-wise attention maps generated by CMA model and general self-attention respectively. It is observed that the CMA model concentrates more

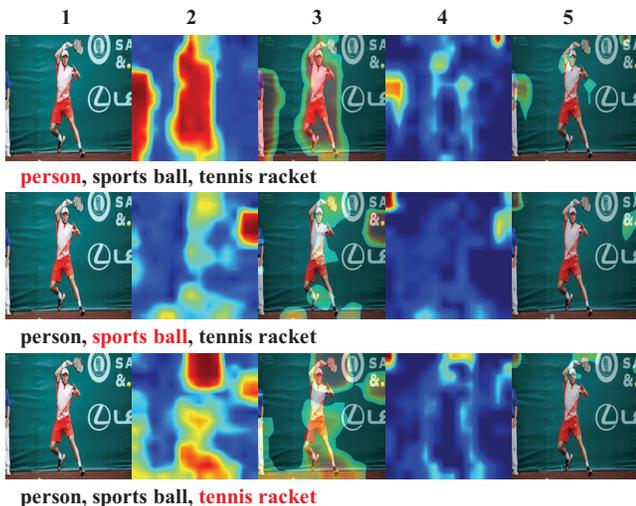


Figure 4: The visualization of attention maps. The first column: original image, second and fourth column: attention maps of MS-CMA and self-attention respectively, third and fifth column: attention maps projected on the original image of MS-CMA and self-attention respectively.

on semantic regions and has stronger response than general self-attention, thus it is capable of exploiting more discriminative and meaningful information. Besides, our CMA mechanism has the ability of capturing the spatial semantic dependencies, especially for the indiscernible or small objects occur in the image, e.g. attention of sports ball also pays attention to tennis racket due to their semantic similarity. It’s quite helpful because these objects need richer contextual cues to help recognition.

## 4.2 Multi-label Video Classification

**Implementation Details.** For the training of ASGE, we apply optimization relaxation and set  $\alpha = 0.1$ . Other settings are same as described in MLIC task. For the training of classification, the initial learning rate is 0.0002 and decay each  $2 \times 10^6$  samples with momentum 0.8. The hyperparameter  $\beta$  in the Eq.12 is 0. The optimizer is SGD with momentum 0.9. The batch size is 256. In this task, we only implemented a single scale CMA model.

**Evaluation Metrics.** In the MLVC task, we use several metrics to evaluate our model, including Global Average Precision (GAP)(Shin et al. 2018), Average Hit Rate (Avg Hit@1), Precision at Equal Recall Rate (PERR) and Mean Average Precision (mAP)(Abu-El-Haija et al. 2016).

**Results on YouTube-8M Segments Dataset.** In the MLVC task, we verify the effectiveness of our model on the YouTube-8M Segments dataset, which is an extension of the YouTube-8M dataset(Abu-El-Haija et al. 2016).

In our experiment, we only use frame-level image features, while the state-of-the-art methods use additional audio features and most are built on model ensemble, which is unfair to compare with. For this reason, we compare our CMA

Methods	Avg Hit@1	Avg PERR	mAP	GAP
Self-attention-FC	85.1	79.1	49.6	81.8
Self-attention-SNet	86.1	80.2	53.2	83.3
<b>CMA-FC</b>	85.4	79.5	50.7	82.3
<b>CMA-SNet</b>	<b>86.7</b>	<b>81.0</b>	<b>55.8</b>	<b>84.1</b>

Table 4: The comparison between self-attention model and ours CMA model on the YouTube-8M Segments dataset.

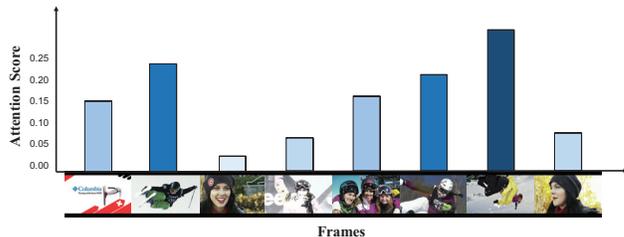


Figure 5: The attention scores for each frame. The label is *skiing*.

model with general self-attention model to validate the effectiveness in MLVC task. Besides, in order to explore the impact of backbone network on CMA mechanism, we implement the SNet-based and FC-based (two fully connected layers) models. The quantitative results are shown in the Table 4. It can be found that all of our metrics are better than the self-attention model. It seems that the improvements compared with self-attention model are not so significant as that of MLIC, but considering the input of our model is fixed pre-extracted features and the models only differ in attention mechanism, the performance gains are quite remarkable.

**Visualization and Analysis.** We also present the visualization results of attention scores for a video in Fig.5. The label of the video is *skiing*, and our CMA model pays more attention to *skiing*-related frames while partly ignores the redundant frames, suggesting that our attention mechanism is capable of locating attentional frames and demonstrating the effectiveness of our model more intuitively.

## 5 Conclusion

In this paper, we propose a novel cross-modality attention mechanism with semantic graph embedding for both MLIC and MLVC task. The proposed method can effectively discover semantic location with rich discriminative features and capture the spatial or temporal dependencies between labels. The extensive evaluations on two MLIC datasets MS-COCO and NUS-WIDE show our method outperforms state-of-the-arts. In addition, we conducted experiments on MLVC dataset YouTube-8M Segments and achieve excellent performance, which validate the strong generalization of our method.

## 6 Acknowledgement

This work was supported by the National Natural Science Foundation of China(61671397). We thank all anonymous reviewers for their constructive comments.

## References

- Abu-El-Haija, S.; Kothari, N.; Lee, J.; Natsev, P.; Toderici, G.; Varadarajan, B.; and Vijayanarasimhan, S. 2016. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*.
- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. Netvlad: Cnn architecture for weakly supervised place recognition. In *CVPR*.
- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*.
- Campos, V.; Jou, B.; Giró-i Nieto, X.; Torres, J.; and Chang, S.-F. 2017. Skip rnn: Learning to skip state updates in recurrent neural networks. *arXiv preprint arXiv:1708.06834*.
- Cao, S.; Lu, W.; and Xu, Q. 2015. Grarep: Learning graph representations with global structural information. In *CIKM*. ACM.
- Chen, S.-F.; Chen, Y.-C.; Yeh, C.-K.; and Wang, Y.-C. F. 2018. Order-free rnn with visual attention for multi-label classification. In *AAAI*.
- Chen, Z.-M.; Wei, X.-S.; Wang, P.; and Guo, Y. 2019. Multi-label image recognition with graph convolutional networks. In *CVPR*.
- 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 *ICIVR*. ACM.
- Gan, C.; Wang, N.; Yang, Y.; Yeung, D.-Y.; and Hauptmann, A. G. 2015. Devnet: A deep event network for multimedia event detection and evidence recounting. In *CVPR*, 2568–2577.
- Ge, W.; Yang, S.; and Yu, Y. 2018. Multi-evidence filtering and fusion for multi-label classification, object detection and semantic segmentation based on weakly supervised learning. In *CVPR*.
- Hao, Y.; Zhou, J. T.; Yu, Z.; Gao, B. B.; Wu, J.; and Cai, J. 2016. Exploit bounding box annotations for multi-label object recognition. In *CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR*.
- Kahou, S. E.; Bouthillier, X.; Lamblin, P.; Gulcehre, C.; Michalski, V.; Konda, K.; Jean, S.; Froumenty, P.; Dauphin, Y.; Boulanger-Lewandowski, N.; et al. 2016. Emonets: Multimodal deep learning approaches for emotion recognition in video. *Journal on Multimodal User Interfaces*.
- Lee, C.-W.; Fang, W.; Yeh, C.-K.; and Frank Wang, Y.-C. 2018. Multi-label zero-shot learning with structured knowledge graphs. In *CVPR*.
- Li, Q.; Qiao, M.; Bian, W.; and Tao, D. 2016. Conditional graphical lasso for multi-label image classification. In *CVPR*.
- Li, Y.; Song, Y.; and Luo, J. 2017. Improving pairwise ranking for multi-label image classification. In *CVPR*.
- Li, X.; Zhao, F.; and Guo, Y. 2014. Multi-label image classification with a probabilistic label enhancement model. In *UAI*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*. Springer.
- Liu, F.; Xiang, T.; Hospedales, T. M.; Yang, W.; and Sun, C. 2017. Semantic regularisation for recurrent image annotation. In *CVPR*.
- Liu, Y.; Sheng, L.; Shao, J.; Yan, J.; Xiang, S.; and Pan, C. 2018. Multi-label image classification via knowledge distillation from weakly-supervised detection. In *MM*. ACM.
- Long, X.; Gan, C.; De Melo, G.; Liu, X.; Li, Y.; Li, F.; and Wen, S. 2018a. Multimodal keyless attention fusion for video classification. In *AAAI*.
- Long, X.; Gan, C.; De Melo, G.; Wu, J.; Liu, X.; and Wen, S. 2018b. Attention clusters: Purely attention based local feature integration for video classification. In *CVPR*.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *KDD*. ACM.
- Read, J.; Pfahringer, B.; Holmes, G.; and Frank, E. 2011. Classifier chains for multi-label classification. *Machine learning*.
- Romero, A.; Ballas, N.; Kahou, S. E.; Chassang, A.; Gatta, C.; and Bengio, Y. 2014. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*.
- Shin, K.; Jeon, J.; Lee, S.; Lim, B.; Jeong, M.; and Nang, J. 2018. Approach for video classification with multi-label on youtube-8m dataset. In *ECCV*.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *CVPR*.
- 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 *CVPR*.
- Wang, Z.; Chen, T.; Li, G.; Xu, R.; and Lin, L. 2017. Multi-label image recognition by recurrently discovering attentional regions. In *ICCV*.
- Wang, D.; Cui, P.; and Zhu, W. 2016. Structural deep network embedding. In *KDD*. ACM.
- Wei, Y.; Xia, W.; Huang, J.; Ni, B.; Dong, J.; Zhao, Y.; and Yan, S. 2014. Cnn: Single-label to multi-label. *arXiv preprint arXiv:1406.5726*.
- Wu, J.; Ma, L.; and Hu, X. 2017. Delving deeper into convolutional neural networks for camera relocalization. In *ICRA*. IEEE.
- Xu, Z.; Yang, Y.; and Hauptmann, A. G. 2015. A discriminative cnn video representation for event detection. In *CVPR*.
- Zagoruyko, S., and Komodakis, N. 2016. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*.
- Zhang, J.; Wu, Q.; Shen, C.; Zhang, J.; and Lu, J. 2018. Multi-label image classification with regional latent semantic dependencies. *IEEE Transactions on Multimedia*.
- Zhu, F.; Li, H.; Ouyang, W.; Yu, N.; and Wang, X. 2017. Learning spatial regularization with image-level supervisions for multi-label image classification. In *CVPR*.
- Zhu, D.; Cui, P.; Wang, D.; and Zhu, W. 2018. Deep variational network embedding in wasserstein space. In *KDD*. ACM.