

# Semantic-Aware Dual Contrastive Learning for Multi-Label Image Classification

Leilei Ma<sup>a,d,†</sup>, Dengdi Sun<sup>b,d,†</sup>, Lei Wang<sup>c</sup>, Haifeng Zhao<sup>a,d,e,‡</sup> and Bin Luo<sup>a,d</sup>

<sup>a</sup>School of Computer Science and Technology, Anhui University, China

<sup>b</sup>School of Artificial Intelligence, Anhui University, China

<sup>c</sup>School of Computer Science and Engineering, Nanjing University of Science and Technology, China

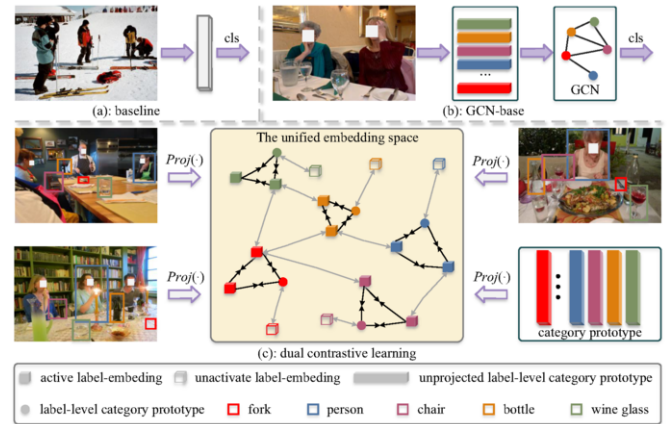
<sup>d</sup>Anhui Provincial key Laboratory of Multimodal Cognitive Computing, Anhui University, China

<sup>e</sup>Institute of Artificial Intelligence, Hefei Comprehensive National Science Center, China

**Abstract.** Extracting image semantics effectively and assigning corresponding labels to multiple objects or attributes for natural images is challenging due to the complex scene contents and confusing label dependencies. Recent works have focused on modeling label relationships with graph and understanding object regions using class activation maps (CAM). However, these methods ignore the complex intra- and inter-category relationships among specific semantic features, and CAM is prone to generate noisy information. To this end, we propose a novel semantic-aware dual contrastive learning framework that incorporates sample-to-sample contrastive learning (SSCL) as well as prototype-to-sample contrastive learning (PSCL). Specifically, we leverage semantic-aware representation learning to extract category-related local discriminative features and construct category prototypes. Then based on SSCL, label-level visual representations of the same category are aggregated together, and features belonging to distinct categories are separated. Meanwhile, we construct a novel PSCL module to narrow the distance between positive samples and category prototypes and push negative samples away from the corresponding category prototypes. Finally, the discriminative label-level features related to the image content are accurately captured by the joint training of the above three parts. Experiments on five challenging large-scale public datasets demonstrate that our proposed method is effective and outperforms the state-of-the-art methods. Code and supplementary materials are released on <https://github.com/yu-gi-oh-leilei/SADCL>.

## 1 Introduction

Multi-label image classification (MLIC) aims to assigning multiple labels to objects or attributes present in a natural image. As a fundamental task in computer vision, it is an essential component in many applications, such as attribute recognition [32], weakly supervised semantic segmentation [22] and automatic image annotation [13]. In general, distinct from single-label classification, *the MLIC task faces two main challenges*: i) complex intra- and inter-category relationships which are hard to be modeled, and ii) various object scales, appearances, and layouts which make it challenging to extract image semantic information effectively.



**Figure 1:** Motivation of the proposed dual contrastive learning. (a) The baseline methods obtain a feature representation of multiple-category blends. (b) GCN-based methods can easily cause overfitting. (c) With the intervention of category prototypes and contrastive learning, our method can obtain discriminative features.

For the second challenge, region- and attention-based methods show their ability to discover objects with various scales and discontinuous regions. However, region-based approaches still have some drawbacks, *e.g.*, high computational overhead [29] and a large number of noisy region proposals [27]. For example, You *et al.* design cross-modal attention modules with cosine similarity to learning category-related regional features [34]. Zhao *et al.* develop a CAM-based module to generate category-specific activation maps which are further leveraged to convert cross-scale semantic feature maps into semantic-aware features [35]. Nevertheless, the aforementioned methods have the following limitations: i) the localized semantic region or object region lacks discrimination, ii) only inter-category relationships (intra-image) are considered, and intra-category relationships (cross-image) are ignored.

To effectively address these limitations as well as the major challenges in multi-label image classification, we propose semantic-aware dual contrastive learning (SADCL) strategy, which is beyond image-level contrastive learning [14] in building inter-category and intra-category discriminative correlation jointly. **To address the first challenge**, our SADCL constructs contrastive representation learning from two aspects, *i.e.*, sample-to-sample and prototype-to-

\* † Equal Contribution

‡ Corresponding Author. Email: senith@163.com.

sample, as shown in Figure 1 (c). For the former, the semantic-aware representations, or label-level visual representations, from both intra- and cross-images are contrasted. Since multiple categories may occur concurrently in a multi-label image, conventional contrastive learning, which only considers intra-category contrastive representations from cross-image, is insufficient. Hence, we also add inter-category representations from single image for sample-to-sample contrastive learning, where modeling label correlations is involved. For the latter, we introduce visual prototypes and produce sample pairs from label-level visual representations and their corresponding prototypes. The visual prototype of one category preserves representative information of this category and can be viewed as a faithful category agent. The prototype-to-sample contrastive learning contributes to enlarging/decreasing the similarity between positive/negative label-level visual representations and their prototypes for each image. Note that we use the same project head for both contrastive learning, so the learned representations are in a unified embedding space. In addition, **to handle the second challenge**, we model the context relationship among multiple objects and scenes at the front end of the framework, and generate an initial label-level visual representation with abundant semantic information through transformer autoencoder with multi-head attention mechanism. Then, the above dual contrastive learning is integrated to further optimize the label-level features in the unified embedding space, so as to obtain the more discriminative semantic-aware representations of an image.

Overall, in our method, discriminative label-level visual representations have the following characteristics: i) they are derived from local discriminative regions of objects, ii) in the unified embedding space, negative label-level visual representations are more discriminative against associated category prototypes than positive ones, and iii) label-level visual representations are largely distinct between each category-pair with low co-occurrence, while similar between that with high co-occurrence.

Our contributions can be summarized in three-fold:

- We propose a novel Semantic-Aware Dual Contrastive Learning framework named SADCL for multi-label image classification, effectively learning more discriminative feature representation.
- Compared with class activation mapping (CAM), we leverage Semantic-Aware Representation Learning to accurately and easily locate the label-related image regions.
- Experiments on five challenging large-scale public datasets (MS-COCO, PASCAL VOC 2007&2012, NUS-WIDE, and Visual Genome) show that our proposed method is effective and outperforms the state-of-the-art methods.

## 2 Related Works

**Multi-Label Image Classification.** With the development of computer vision, multi-label image classification (MLIC) task has received extensive attention and made remarkable progress. The related researches could be roughly divided into two aspects *i.e.*, relation-based methods and region/attention-based methods.

Since the relationship of multiple objects co-occurring in a multi-label image can be meaningful for classification, modeling label correlation has become a hot topic in MLIC. Pioneer works [25, 26] utilize RNN or LSTM to transform labels and an image into a joint embedding space and predict labels in a pre-defined orderly way. Despite this, RNN-based methods only model label-local relations while ignoring label-global relations. To address this limitation, GCN-based methods propagate information between nodes over a graph and explore labels-wise relationships. For example, Li *et*

al. [18] design an adaptive label graph learning module with sparse correlation constraints to reduce the hassle of hand-crafted graphs. To overcome the dependence of GCN on learning only pair-wise labels, Wu *et al.* [28] introduce adaptive hypergraph neural network to model the higher-order semantic relationships among labels automatically. On the other hand, region and attention-based methods have been handling various object scales and appearances in MLIC tasks. Many existing works [29] adopt object detection methods to generate a set of semantic-aware instances and construct semantic label graphs. Nevertheless, object detection methods require the pre-trained detector with additional bounding box annotations. Exploiting attention mechanism to extract salient object regions is another popular fashion. For example, Ye *et al.* [33] design a semantic attention module to generate category-specific activation maps, and then obtain content-aware category representations. Zhao *et al.* [35] propose a cross-attention module to suppress the noise between different scales and enhance the structural information of small objects. Liu *et al.* [21] make use of label embeddings as queries to learn class-specific representations via cross-modal attention.

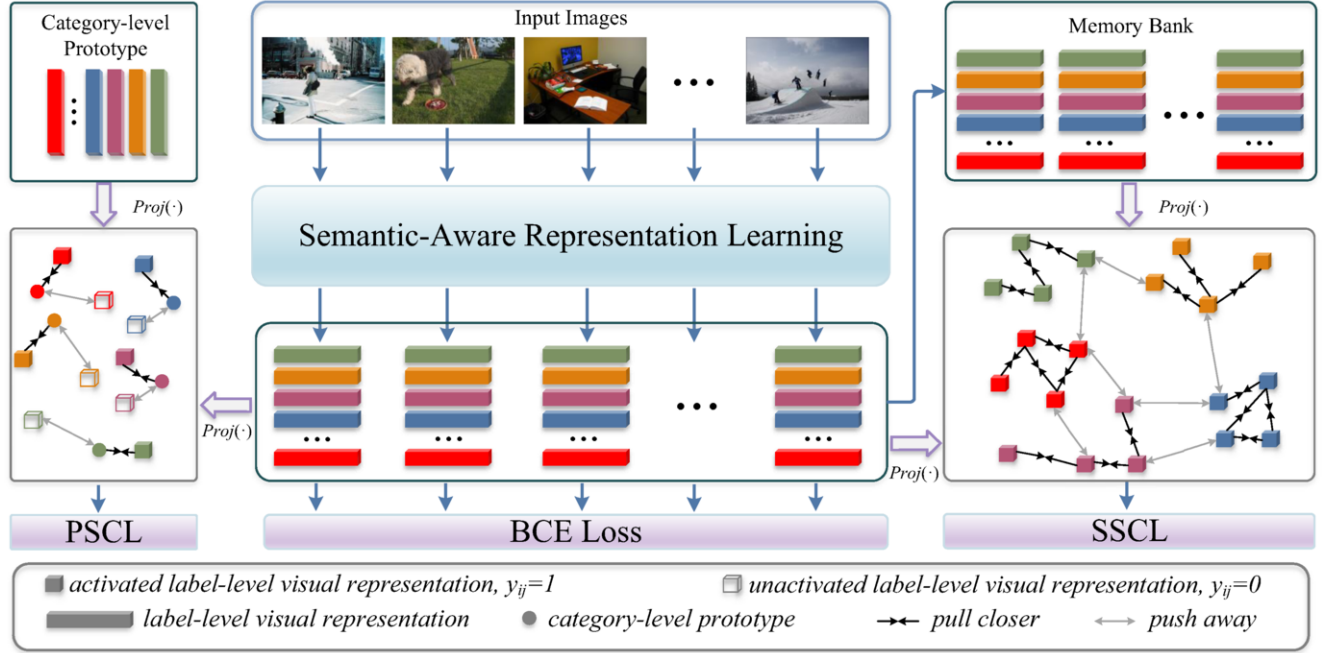
**Contrastive learning.** Contrastive learning is an effective method to strengthen the discrimination of the learned representations, which aims to bring the samples in positive sample pairs closer together in a unified embedding space while pushing samples in negative pairs away. For example, He *et al.* [11] build a dynamic dictionary with a queue and a moving average encoder for unsupervised contrastive learning. Khosla *et al.* [14] leverage label information for supervised contrastive learning, which sets samples with the same label as positive pairs while samples with different labels as negative pairs. In addition, Li *et al.* [17] proposed prototypical contrastive learning for unsupervised presentation learning, where prototypes represent class-agnostic semantic information, but we propose prototypes for class-specific semantic representations. Although contrastive learning for MLIC has been explored preliminarily by Dao *et al.* [8], it fails to model label correlation. Moreover, two-time augmentation of an image makes it inefficient.

## 3 Methodology

### 3.1 Overview and Preliminary

In an MLIC task, given a minibatch of input images  $\mathcal{D} = \{(I_i, Y_i)\}_{i=1}^N$  with  $L$  categories, where  $N$  is the batch size, our goal is to build a visual model  $F(\cdot)$  to predict the label  $Y_i$  of the given image  $I_i$ . If an image  $I_i \in \mathcal{D}$  contains the  $j$ -th category, the associated label  $y_{ij} \in Y_i$  will be 1, and vice versa.

For this purpose, we propose a Semantic-Aware Dual Contrast Learning (SADCL) framework that learns more discriminative representations for MLIC task. As shown in Figure 2, the architecture of the proposed SADCL consists of three components: Semantic-Aware Representation Learning (SARL), Sample-to-Sample Contrastive Learning (SSCL), and Prototype-to-Sample Contrastive Learning (PSCL). SARL decomposes the image into  $L$  label-level visual representations (features), and then SSCL and PSCL optimize the distribution of label-level representations from the perspective of samples and category prototype respectively, so that in a unified embedded space the homogeneous and activated features tend to aggregate, while heterogeneous or unactivated features are pushed away. Finally, the optimized label-level representation is fed into a classifier to predict whether each category is present or not. In the inferring stage, no more contrastive learning is needed, and we use SARL to directly obtain label-level visual representation for classification.



**Figure 2:** The overview of the SADCL framework. **Note that a solid square represents an active state, and a hollow square represents an inactive state.** Meanwhile, different colors represent different categories.  $\text{Proj}(\cdot)$  is a projection network.

### 3.2 Semantic-Aware Representation Learning

To accurately extract semantic information from complex image content, following [8, 21], we first model the contextual spatial layout among multiple objects and scene to generate feature maps with abundant contextual semantic information. Specifically, the initial spatial feature maps  $\mathcal{F} \in \mathbb{R}^{H_0 \times W_0 \times d}$  of the input image are extracted through CNN, where  $H_0$ ,  $W_0$  and  $d$  are the length, width, and channel number, and reshaped as a sequence feature form:  $\mathcal{F}^0 \in \mathbb{R}^{H_0 W_0 \times d}$ . Then the sequential features  $\mathcal{F}^0$  and the corresponding positional embedding  $PE$  are input to a transformer-encoder, as shown on the left half of Figure 3. In this way, the output features with rich contextual semantic information will be generated as follows:

$$\mathcal{F}^i = \text{FFN}(\text{SelfAttBlock}(\tilde{\mathcal{F}}^{i-1}, PE)), \quad (1)$$

where multi-head self-attention (SelfAttBlock) and feed-forward neural network (FFN) are the main components of the transformer-encoder. Eq. 1 is a layerwise iterative form, and here we can simply take the output features to be  $\mathcal{F}^1$ .

After that, as shown on the right half of Figure 3, we utilize the learnable category semantic embeddings as queries  $Q^{in} \in \mathbb{R}^{C \times d}$ , and the spatial features  $\mathcal{F}^i$  as keys and values, and perform a transformer-decoder to generate the label-level visual representations as follows.

$$\mathcal{Q}^i = \text{FFN}(\text{CrossAttBlock}(\tilde{\mathcal{F}}^{i-1}, \tilde{Q}^{i-1}, PE)), \quad (2)$$

where multi-head cross-attention (CrossAttBlock) locates category-related discriminative regions, and different heads focus on information from different parts of the same object. Consequently, the regions in the feature maps with high correlation to  $\mathcal{Q}^{i-1}$  are aggregated in  $\mathcal{Q}^i$  and updated layer-by-layer. In this way, the label-related regions can be located as accurately as possible to extract semantic-aware visual representations. For example, if an image contains a *cat*, the region in the image from the *cat* will be highlighted and have a

higher correlation with the corresponding category semantic embedding.

### 3.3 Projection Network for Contrastive Learning

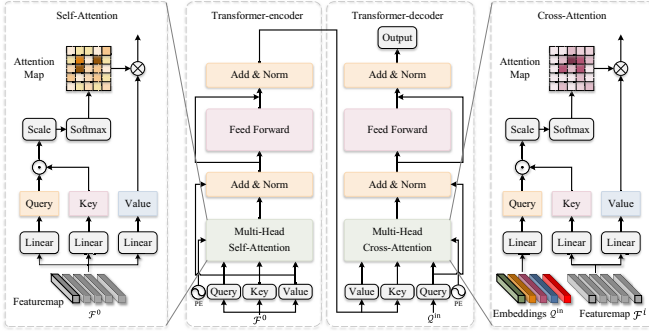
After capturing the initial label-level visual representations, perform a projection transformation and embed the learned features into a unified vector space for dual contrastive learning. As mentioned above, given an image  $I_i$ , the label-level features  $Q^{out}$  is generated by backbone and semantic-aware representation learning. Then,  $Q^{out}$  is mapped to a unified embedding space using a projection network, which consists of two linear layers and an activation function:

$$x_i = \text{Proj}(Q^{out}) \in \mathbb{R}^{L \times d}, \quad (3)$$

where  $x_i$  is the projected label-level visual representation for the image  $I_i$ . Since contrastive learning is conducted in the projected space, we can also refer the unified embedding space to the projected space.

### 3.4 Sample-to-Sample Contrastive Learning

In the single-label image classification task, most previous contrastive learning methods [11, 14] are based on image-level contrastive losses. These methods define positive and negative sample pairs as follows: in a mini-batch of the input images, the image-level representations with different augmentations belonging to the same category are defined as positive sample pairs, and the image-level representations from other categories or images are defined as negative sample pairs. However, for multi-label image classification, an image usually contains multiple categories, making it challenging to define positive or negative sample pairs at the image level. Instead of using image-level representations like existing works on contrastive learning, we propose sample-to-sample contrastive learning, which defines positive and negative sample pairs at the label level. In the minibatch of  $N$  input images, the projected label-level



**Figure 3:** Illustrations of Transformer-encoder and Transformer-decoder.

visual representations from the memory bank and SADCL are aggregated into a set  $\mathcal{X} = \{(x_{ij}, y_{ij}) | i \in \{1, \dots, 2N\}; j \in \{1, \dots, L\}\}$ , where  $x_{ij} \in \mathbb{R}^d$  is the label-level visual representation about the  $j$ -th category in the  $i$ -th image,  $y_{ij} \in \{0, 1\}$  is the ground-truth associated with  $x_{ij}$ . Hereafter, we view  $x_{ij}$  as an activated label-level visual representation, when  $y_{ij} = 1$ , vice versa. After that, we define positive and negative sample pairs in the activated label-level visual representation set  $\mathcal{I} = \{x_{ij} \in \mathcal{X} | y_{ij} = 1\}$ . Here, the  $x_{ij} \in \mathcal{I}$  is called the anchor, and the  $x_{ij(p)}$  is referred to as the positive. Note that the positive set is formulated in the following form  $x_{ij(p)} \in \mathcal{P}_{ss}(i, j) = \{x_{kj} \in \mathcal{A}(i, j) | y_{kj} = y_{ij} = 1\}$ , where  $\mathcal{A}(i, j) = \mathcal{I} \setminus x_{ij}$  is the set including  $\mathcal{I}$  and excluding  $x_{ij}$ . In summary, the sample-to-sample contrastive loss for the anchor  $x_{ij}$  are defined as follows:

$$\begin{aligned} \mathcal{L}_{S2S}^{ij} &= \sum_{x_{ij(p)} \in \mathcal{P}_{ss}(i, j)} \log \frac{\exp(x_{ij} \cdot x_{ij(p)} / \tau)}{\sum_{x_a \in \mathcal{A}(i, j)} \exp(x_{ij} \cdot x_a / \tau)}, \\ \mathcal{L}_{S2S} &= \sum_{x_{ij} \in \mathcal{I}} \frac{-1}{|\mathcal{P}_{ss}(i, j)|} \mathcal{L}_{S2S}^{ij}, \end{aligned} \quad (4)$$

where  $\tau \in \mathcal{R}^+$  denotes a temperature parameter.

In addition, we propose a memory bank module because an image needs to be augmented twice in the conventional contrastive learning approach, which will double the image encoding calculation. The memory bank stores activated label-level visual representations ( $y_{i,j} = 1$ ) from each training iteration, replacing the original representations.

### 3.5 Prototype-to-Sample Contrastive Learning

SSCL only considers using activated label-level visual representations ( $y_{ij} = 1$ ) to perform contrastive learning. Furthermore, an image contains only a small number of categories, resulting in many label-level visual representations that are not active ( $y_{ij} = 0$ ). This leads to inactive label-level visual representations that do not participate in contrastive learning, thereby restricting the model from obtaining more discriminative representations. To alleviate this problem, we design a prototype-based contrastive learning loss to mine this unactivated label-level visual representation information entirely. As shown in Figure 2, in the projected space, the activated label-level visual representations are as close as possible to the corresponding category prototypes. In contrast, inactive label-level visual representations are as far away as possible from the center of the related category prototypes.

In line with this goal, we propose a category-level prototype module, which can be regarded as the center of the category. The label-level visual representations produced during model training are averaged to obtain category-level prototypes  $\mathcal{C}^{in}$ . To ensure that the label-level visual representation is reliable, we propose a screening scheme:  $\mathcal{Q}^{out}$  is the activation state ( $y_{ij} = 1$ ), and the corresponding prediction score is greater than a fixed threshold  $\epsilon$ . The process to obtain the category-level prototype is formulated as follows:

$$\begin{aligned} \mathcal{C}_j^{in} &= \frac{1}{N_j} \sum_{i=1}^N \mathcal{Q}_{ij}^{out} \cdot \mathbb{I}(\text{pred}(\mathcal{Q}_{ij}^{out}) \geq \epsilon) \cdot \mathbb{I}(y_{ij} = 1), \\ \mathcal{C}^{in} &= [\mathcal{C}_1^{in}, \dots, \mathcal{C}_j^{in}, \dots, \mathcal{C}_L^{in}]^T \in \mathbb{R}^{L \times d}, \end{aligned} \quad (5)$$

where  $N_j$  is the number of label-level visual representations  $\mathcal{Q}_{ij}^{out}$  related to category  $j$ . As a result of mapping  $\mathcal{C}^{in}$  to the projection space,  $\mathcal{C}^{out} \in \mathbb{R}^{L \times d}$  is obtained for contrastive learning by using the projection network. For a specific category  $j$ , if we define  $p_j$  as an anchor, the label-level visual representation for the category  $j$ : the activated  $x_{ij(p)}$  as the set of positive set  $\mathcal{P}_{sp}(i, j)$  and the inactivated  $x_{ij}$  as the set of negative set  $\mathcal{N}_{sp}(i, j)$ .

Based on these analyses, the prototype-to-sample contrastive loss is designed as:

$$\mathcal{L}_{P2S} = - \sum_{j \in \mathcal{C}} \log \frac{\sum_{x_{ij(p)} \in \mathcal{P}_{sp}(i, j)} \exp(c_j \cdot x_{ij(p)} / \tau)}{\sum_{x_{ij} \in \mathcal{P}_{sp}(i, j) \cup \mathcal{N}_{sp}(i, j)} \exp(c_j \cdot x_{ij} / \tau)} \quad (6)$$

### 3.6 Classification and Optimizing

With the label-level visual representations  $\mathcal{Q}^{out}$  obtained by SARL, we treat each  $\mathcal{Q}^{out}$  prediction as a binary classification task and design a group of binary classifiers, consisting of a linear transformation layer with sum operation and a nonlinear activation function:

$$s_l = \sigma \left( \text{sum} \left( W_l^T \mathcal{Q}_l^{out} + b_l \right) \right), \quad (7)$$

where  $W_l^T \in \mathbb{R}^d$  is a learnable weight for category  $l$ ,  $b_l \in \mathbb{R}^1$  is a bias,  $\sigma$  is sigmoid activation, and  $s_l$  denotes the prediction score.

Following conventional multi-label image classification works, we adopt BCE loss as the classification loss:

$$\mathcal{L}_{BCE} = \sum_{j=1}^L y_{ij} \log s_{ij} + (1 - y_{ij}) \log(1 - s_{ij}), \quad (8)$$

where  $s_{ij}$  is the score for predicting  $x_{ij}$ , and  $y_{ij}$  is the corresponding ground-truth label. As a final step, classification loss, sample-to-sample contrastive loss and prototype-to-sample contrastive learning jointly optimize our model:

$$\mathcal{L} = \mathcal{L}_{S2S} + \mathcal{L}_{P2S} + \mathcal{L}_{BCE}. \quad (9)$$

## 4 Experiments

### 4.1 Experimental Setup

**Datasets.** To evaluate the proposed method, we used the five most popular benchmark datasets: MS-COCO [19], NUS-WIDE [7], Visual Genome [15], PASCAL VOC 2007 [9], and PASCAL VOC 2012 [9]. All data collections are natural image collections containing categories commonly found in natural scenes. The details of each dataset are shown in Table 2.



Methods	$(R_{train}, R_{test})$	mAP	All						Top 3					
			CP	CR	CF1	OP	OR	OF1	CP	CR	CF1	OP	OR	OF1
CNN-RNN [25]	(-, -)	61.2	-	-	-	-	-	-	66.0	55.6	60.4	69.2	66.4	67.8
RNN-Att [26]	(-, -)	-	-	-	-	-	-	-	79.1	58.7	67.4	84.0	63.0	72.0
ResNet101* [12]	(448, 448)	81.5	82.1	71.2	76.0	84.6	75.4	79.7	85.9	62.9	71.6	89.6	66.1	76.1
MLGCN [5]	(448, 448)	83.0	85.1	72.0	78.0	85.8	75.4	80.3	89.2	64.1	74.6	90.5	66.5	76.7
MS-CMA [34]	(448, 448)	83.8	82.9	74.4	78.4	84.4	77.9	81.0	88.2	65.0	74.9	90.2	67.4	77.1
P-GCN [6]	(448, 448)	83.2	84.9	72.7	78.3	85.0	76.4	80.5	89.2	64.3	74.8	90.0	66.8	76.7
GM-MLIC [29]	(448, 448)	84.3	87.3	70.8	78.3	88.6	74.8	80.6	90.6	67.3	74.9	94.0	69.8	77.8
MCAR [10]	(448, 448)	83.8	85.0	72.1	78.0	88.0	73.9	80.3	88.1	65.5	75.1	91.0	66.3	76.7
TDRG [35]	(448, 448)	84.6	86.0	73.1	79.0	86.6	76.4	81.2	89.9	64.4	75.0	91.2	67.0	77.2
CCD-R101 [20]	(448, 448)	84.0	87.2	70.9	77.3	88.8	74.6	81.1	89.7	63.9	72.9	92.0	66.5	77.2
MulCon [8]	(448, 448)	84.9	84.0	74.8	79.2	85.6	78.0	81.6	87.8	65.9	75.3	90.5	67.9	77.6
Query2Label [21]	(448, 448)	84.9	84.8	74.5	79.3	86.6	76.9	81.5	78.0	69.1	73.3	80.7	70.8	75.4
CPSD [30]	(448, 448)	84.9	88.4	71.7	79.2	89.3	74.8	81.4	-	-	-	-	-	-
Ours(SADCL)	(448, 448)	85.6	84.6	76.0	79.8	86.0	78.5	82.1	88.9	66.6	74.9	91.0	68.3	78.0
ADDGCN [33]	(448, 576)	85.2	84.7	75.9	80.1	84.9	79.4	82.0	88.8	66.2	75.8	90.3	68.5	77.9
SSGRL [2]	(576, 576)	83.8	89.9	68.5	76.8	91.3	70.8	79.7	91.9	62.5	72.7	93.8	64.1	76.2
AdaHGNN [28]	(576, 576)	85.0	-	-	79.9	-	-	81.8	-	-	75.5	-	-	77.6
C-Tran [16]	(576, 576)	85.1	86.3	74.3	79.9	87.7	76.5	81.7	90.1	65.7	76.0	92.1	71.4	77.6
TDRG [35]	(576, 576)	86.0	87.0	74.7	80.4	87.5	77.9	82.4	90.7	65.6	76.2	91.9	68.0	78.1
CCD-R101 [20]	(576, 576)	85.3	88.3	73.1	80.2	88.8	76.3	82.1	91.0	65.2	76.0	92.3	67.3	77.9
Ours(SADCL)	(448, 576)	86.8	86.4	77.0	81.1	87.7	79.1	83.2	90.0	67.4	75.7	92.0	68.7	78.7

**Table 1:** Comparisons with state-of-the-art methods on the MS-COCO dataset.\* indicates the reproduced results of our implementation. All metrics are in %.

**Evaluation Metrics.** To evaluate our method, we adopt the average precision (AP) per category and mean of average precision (mAP) across all categories for evaluation. As the three datasets (MS-COCO, NUS-WIDE, Visual Genome) contain so many categories, it is impossible to present the AP of each category, so in addition to mAP, we also show the following evaluation metrics: overall precision (OP), recall (OR), F1-measure (OF1) and per-category precision (CP), recall (CR), and F1-measure (CF1), computed over all prediction scores and top-3 highest prediction scores. All metrics are in %.

**Implementation Details.** For fair comparisons, ResNet-101 [12] pre-trained on ImageNet-1k is adopted as our backbone. To accelerate model convergence, we use AdamW as the optimizer, and OneCycleLR is used as the learning rate scheduler, with the maximum learning rate being  $5e-5$  under a batch size of 64. For data augmentation and regularization, we employ the tricks suggested in [21]: RandAugment, Cutout with a factor of 0.5, weight decay of  $5e-3$ . Our method can converge fast, which achieves optimal results in less than 20 epochs.

## 4.2 Comparison with State-of-The-Art Methods

**Results on MS-COCO.** MS-COCO is the most widely used large-scale dataset with challenges for multi-label image classification. It has the following characteristics: common categories, various object scales, and imbalanced classes. As shown in Table 1, compared with state-of-the-art methods, our method substantially outperforms other methods in all evaluation metrics, especially on mAP, CF1, and OF1. Most methods use  $448 \times 448$  image resolution as input during the training and testing stage, but some works change the input image resolution to  $576 \times 576$ . To provide a fair comparison, we report the results only after modifying the input resolution to  $576 \times 576$  during the testing phase.

Dataset	#Train	#Test	#Category	#Label Cardinality
MS-COCO	82783	40504	80	2.9
NUS-WIDE	125449	83898	81	2.4
VOC2007	5011	4952	20	1.5
VOC2012	11540	10991	20	1.4
VG500	82904	10000	500	13.6

**Table 2:** Statistics for the popular multi-label image classification dataset, including image numbers, category numbers, and label cardinality. Note that Visual Genome is referred to as VG500, PASCAL VOC 2007 is called VOC2007, and PASCAL VOC 2012 is named VOC2012.

**Results on PASCAL VOC 2007 and 2012.** PASCAL VOC 2007 and 2012 are used as datasets for object detection and segmentation and evaluated for the MLIC task. Since these two datasets cover 20 common categories, we report AP per category in addition to mAP. As shown in Table 3, in terms of mAP, our proposed method achieves 96.4%, outperforming other methods with a larger margin. In addition, the increase in the number of images trained by VOC2012 further influences the results. As shown in Table 4, our method outperforms existing methods on VOC2012 with a larger margin. Note that our method also outperforms other methods in small object recognition, e.g. *bird*, *cat*, and *plant*.

**Results on NUS-WIDE.** NUS-WIDE is a web image dataset, with all images coming from Flickr. The experimental results is reported in Table 5. Compared with the existing state-of-the-art methods, our method achieves significant performance on five evaluation metrics with 65.9% mAP, 63.0% CF1(All), 75.0% OF1(All), 57.8% CF1 (Top 3), and 70.6% OF1 (Top 3). NUS-WIDE's image quality (resolution) is lower than MS-COCO's, so the recognition performance on NUS-WIDE of all methods has yet to reach the level of MS-COCO.

**Results on Visual Genome.** Visual Genome is a dataset recently used to evaluate multi-label image classification. Following SS-

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mAP
CNN-RNN [25]	96.7	83.1	94.2	92.8	61.2	82.1	89.1	94.2	64.2	83.6	70.0	92.4	91.7	84.2	93.7	59.8	93.2	75.3	99.7	78.6	84.0
ResNet-101* [12]	99.0	98.4	97.5	96.0	81.4	97.3	97.3	97.1	79.6	96.0	88.1	97.5	98.5	95.8	98.8	85.9	97.2	84.6	98.8	92.0	93.8
RNN-Att [26]	98.6	97.4	96.3	96.2	75.2	92.4	96.5	97.1	76.5	92.0	87.7	96.8	97.5	93.8	98.5	81.6	93.7	82.8	98.6	89.3	91.9
SSGRL* [2]	99.5	97.1	97.6	97.8	82.6	94.8	96.7	98.1	78.0	97.0	85.6	97.8	98.3	96.4	98.1	84.9	96.5	79.8	98.4	92.8	93.4
ML-GCN [5]	99.5	98.5	98.6	98.1	80.8	94.6	97.2	98.2	82.3	95.7	86.4	98.2	98.4	96.7	99.0	84.7	96.7	84.3	98.9	93.7	94.0
P-GCN [6]	99.6	98.6	98.4	98.7	81.5	94.8	97.6	98.2	83.1	96.0	87.1	98.3	98.5	96.3	99.1	87.3	95.5	85.4	98.9	93.6	94.3
ADDGCN* [33]	99.8	99.0	98.4	99.0	86.7	98.1	98.5	98.3	85.8	98.3	88.9	98.8	99.0	97.4	99.2	88.3	98.7	90.7	99.5	97.0	96.0
KGGR* [1]	99.3	98.6	97.9	98.4	86.2	97.0	98.0	99.2	82.6	98.3	87.5	99.0	98.9	97.4	99.1	86.9	98.2	84.1	99.0	95.0	95.0
DSDL [36]	99.8	98.7	98.4	97.9	81.9	95.4	97.6	98.3	83.3	95.0	88.6	98.0	97.9	95.8	99.0	86.6	95.9	86.4	98.6	94.4	94.4
GM-MLIC [29]	99.4	98.7	98.5	97.6	86.3	97.1	98.0	99.4	82.5	98.1	87.7	99.2	98.9	97.5	99.3	87.0	98.3	86.5	99.1	94.9	94.7
TDRG [35]	99.9	98.9	98.4	98.7	81.9	95.8	97.8	98.0	85.2	95.6	89.5	98.8	98.6	97.1	99.1	86.2	97.7	87.2	99.1	95.3	95.0
MulCon* [8]	99.8	98.3	99.3	98.6	83.3	98.4	98.0	98.3	85.8	98.3	90.5	99.3	98.9	96.6	98.8	86.3	99.8	87.3	99.8	96.1	95.6
CPCL [30]	99.6	98.6	98.5	98.8	81.9	95.1	97.8	98.2	83.0	95.5	85.5	98.4	98.5	97.0	99.0	86.6	97.0	84.9	99.1	94.3	94.4
SST [3]	99.8	98.6	98.9	85.5	94.7	97.9	98.6	83.0	96.8	85.7	98.8	98.8	98.9	95.7	99.1	85.4	96.2	84.3	99.1	95.0	94.5
Ours(SADCL)*	100.0	99.0	99.5	99.1	88.9	98.8	98.7	99.6	84.2	98.4	90.1	99.4	99.6	99.0	99.3	90.2	99.6	88.9	99.8	95.3	96.4

**Table 3:** Comparisons with state-of-the-art methods on the VOC 2007 dataset. \* indicates the reproduced results of our implementation, and \* indicates the model pre-trained on COCO. All of the inputs are  $448 \times 448$  resolution except SSGRL(576), ADDGCN(512) and KGGR(576).

Methods	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	motor	person	plant	sheep	sofa	train	tv	mAP
VeryDeep [23]	99.1	88.7	95.7	93.9	73.1	92.1	84.8	97.7	79.1	90.7	83.2	97.3	96.2	94.3	96.9	63.4	93.2	74.6	97.3	87.9	89.0
Fev+Lv [31]	98.4	92.8	93.4	90.7	74.9	93.2	90.2	96.1	78.2	89.8	80.6	95.7	96.1	95.3	97.5	73.1	91.2	75.4	97.0	88.2	89.4
HCP [27]	99.1	92.8	97.4	94.4	79.9	93.6	89.8	98.2	78.2	94.9	79.8	97.8	97.0	93.8	96.4	74.3	94.7	71.9	96.7	88.6	90.5
MCAR [10]	99.6	97.1	98.3	96.6	87.0	95.5	94.4	98.8	87.0	96.9	85.0	98.7	98.3	97.3	99.0	83.8	96.8	83.7	98.3	93.5	94.3
SSGRL* [2]	99.7	96.1	97.7	96.5	86.9	95.8	95.0	98.9	88.3	97.6	87.4	99.1	99.2	97.3	99.0	84.8	98.3	85.8	99.2	94.1	94.8
ADDGCN* [33]	99.8	97.1	98.6	96.8	89.4	97.1	96.5	99.3	89.0	97.7	87.5	99.2	99.1	97.7	99.1	86.3	98.8	87.0	99.3	95.4	95.5
KGGR* [1]	99.8	97.3	98.4	97.1	87.9	97.3	96.5	99.3	89.4	97.8	88.7	99.4	99.4	97.9	99.2	86.3	98.8	86.3	99.7	95.2	95.6
DSDL [36]	99.4	95.3	97.6	95.7	83.5	94.8	93.9	98.5	85.7	94.5	83.8	98.4	97.7	95.9	98.5	80.6	95.7	82.3	98.2	93.2	93.2
CCD-R101* [20]	99.8	98.2	98.3	98.0	88.6	97.4	96.9	99.1	90.8	98.9	90.2	99.2	99.6	98.4	99.0	87.7	98.4	88.8	99.7	96.4	96.1
Ours(SADCL)*	99.9	98.7	99.0	98.7	91.2	97.8	97.4	99.6	92.3	98.9	90.3	99.7	99.7	98.4	99.4	89.8	99.1	90.9	99.7	97.2	96.9

**Table 4:** Comparisons with state-of-the-art methods on the VOC 2012 dataset. \* indicates the model pre-trained on COCO. All of the inputs are  $448 \times 448$  resolution except SSGRL(576), ADDGCN(512) and KGGR(576).

Methods	mAP	All		Top 3	
		CF1	OF1	CF1	OF1
CNN-RNN [25]	-	-	-	34.7	55.2
ResNet101* [12]	62.5	59.2	73.8	54.6	69.4
CADM [4]	62.8	60.7	74.1	56.3	70.6
ADDGCN* [33]	63.3	60.3	73.5	56.5	69.3
P-GCN [6]	62.8	60.4	73.4	57.0	69.1
TDRG* [35]	63.5	60.0	73.8	56.1	69.5
SST [3]	63.5	59.6	73.2	55.9	68.8
MulCon [8]	63.9	61.8	74.8	-	-
CCD-R101* [20]	64.2	61.8	74.6	56.7	70.0
Ours(SADCL)	<b>65.9</b>	<b>63.0</b>	<b>75.0</b>	<b>57.8</b>	<b>70.6</b>

**Table 5:** Comparisons with state-of-the-art methods on the NUS-WIDE dataset. \* indicates the reproduced results of our implementation. The input images are resized to  $448 \times 448$  resolution in both the training and testing phases.

GRL [2] setting, we select the 500 most frequently occurring categories to build a new dataset called VG500. Compared to other related datasets, this one presents more challenges due to the large number of categories. In addition to the existing reports, ADDGCN and TDRG are reported using the corresponding official codes. As can be seen from Table 6, we present five important evaluation metrics, including mAP, CF1(All), OP1(All), CF1(Top3), and OP1(Top3). Except for CF1(Top3), our method outperforms other state-of-the-art methods, especially in mAP, obtaining 40.5%.

### 4.3 Ablation Study

In this section, we evaluate the effectiveness of each component of our proposed framework by performing an ablation study on MS-COCO and NUS-WIDE datasets in Table 7. The results show that, on

Methods	Resolution	mAP	All		Top 3	
			CF1	OF1	CF1	OF1
ResNet101*	$448 \times 448$	35.3	29.3	47.0	16.0	27.8
ML-GCN	$448 \times 448$	32.6	27.5	42.8	16.8	27.1
ADDGCN*	$448 \times 448$	37.3	33.2	47.8	19.8	28.4
SSGRL	$576 \times 576$	36.6	-	-	-	-
ADDGCN*	$576 \times 576$	38.2	33.5	47.8	19.9	28.4
KGGR	$576 \times 576$	37.4	32.5	47.2	19.4	28.1
TDRG*	$576 \times 576$	37.7	30.9	48.0	17.5	28.5
C-Tran	$576 \times 576$	38.4	35.2	49.5	<b>20.1</b>	28.7
MulCon	$576 \times 576$	38.5	-	-	-	-
Query2Label	$576 \times 576$	39.5	-	-	-	-
Ours(SADCL)	$576 \times 576$	<b>40.5</b>	<b>41.0</b>	<b>54.0</b>	17.1	<b>29.2</b>

**Table 6:** Comparisons with state-of-the-art methods on the VG500 dataset. \* indicates the reproduced results of our implementation.

both datasets, employing only SARL significantly improves performance compared to the baseline. This improvement is from 81.5% mAP to 85.0% mAP and from 62.5% mAP to 65.3.0% mAP, respectively. This demonstrates that SARL can learn category-related representations by focusing on category-related semantic regions. We observe that SARL\* drops 0.3% mAP on both datasets compared to SARL, which is attributed to the Transformer-Encoder to model long-range dependencies from multiple objects and scene. The collaborative learning of SARL and SSCL improves by 0.4% mAP on MS-COCO and 0.3% mAP on NUS-WIDE compared to SARL, demonstrating that SSCL can learn a better representation of potential features. Combining the SARL and PSCL components, 85.4% mAP and 65.7% mAP were obtained on MS-COCO and NUS-WIDE. This indicates that PSCL can narrow the distance between positive label-level visual representations samples ( $y_{ij} = 1$ ) and category prototypes and increase the distance between negative class samples ( $y_{ij} = 0$ ) and category prototypes. In the end, when

all components are used, the performance of our model is further improved, reaching 85.6% mAP and 65.9% mAP.

Module					MS-COCO			NUS-WIDE		
Baseline	SARL*	SARL	SSCL	PSCL	mAP	OF1	CF1	mAP	OF1	CF1
✓					81.5	79.7	76.0	62.5	73.8	59.2
✓	✓				84.7	81.5	78.7	65.0	74.2	61.0
✓			✓		85.0	81.7	79.3	65.3	74.8	61.7
✓			✓	✓	85.4	81.7	79.6	65.7	74.9	62.7
✓			✓		85.4	82.1	79.8	65.6	75.0	62.4
✓			✓	✓	<b>85.6</b>	<b>82.1</b>	<b>79.8</b>	<b>65.9</b>	<b>75.0</b>	<b>63.0</b>


**Table 7:** Ablation study of different components in MS-COCO and NUS-WIDE Datasets. Baseline mean ResNet101 network with average pooling and fc layer. Noteworthy, when Baseline is combined with additional components, the pool and fc layers in Baseline will be removed. SARL\* denotes SARL after removing Transformer-encoder. All metrics reflect all predicted scores instead of taking the top-3 highest prediction scores.

#### 4.4 Visualization Study

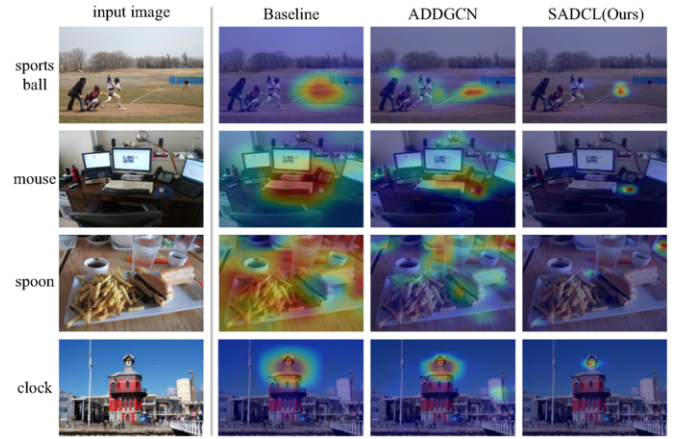
In order to further explore category prototypes, we use t-SNE [24] to visualize them as illustrated in Figure 4. It is clear to show that our category prototypes contain meaningful semantic information. Specifically, the distance between different categories is as large as possible, and paired categories with high co-occurrence have a smaller distance than paired classes with low co-occurrence. Therefore, the correlation between the labels is well modeled by our method, and the category prototypes are discriminative.

To evaluate the performance of our proposed model for localizing discriminative object regions, we visualize category-related attention maps for SARL modules and category-related activation maps for baseline and ADDGCN. The results from Figure 5 show that the baseline method and the ADDGCN method localize ambiguously and easily generate noise. In contrast, our method identifies the most discriminative part of the related object. As shown in the 3th row of Figure 5, from the *spoon* category, we can observe: i) the baseline method fails to localize related regions, ii) the ADDGCN method can focus on *spoon* while generates a lot of noise, and iii) our method can accurately localize discriminative regions for *spoon*.

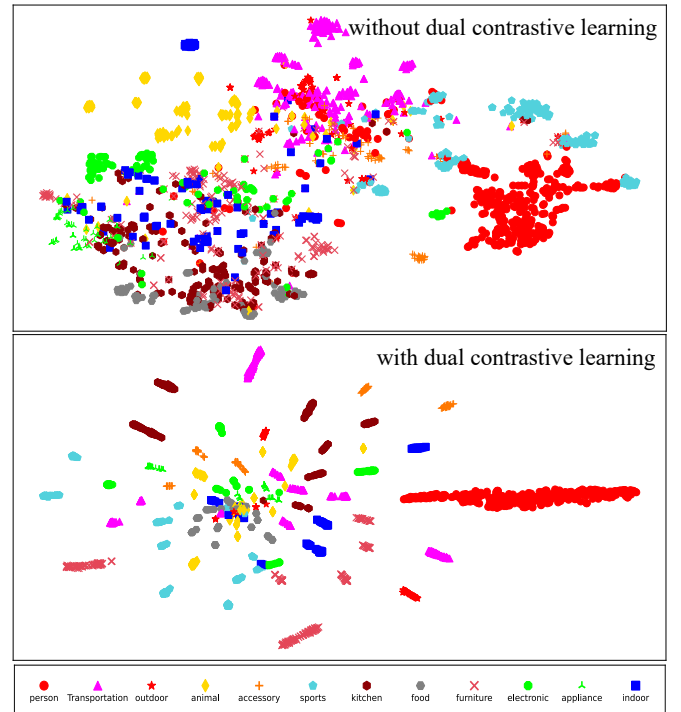
In order to evaluate the quality of the learned representations, we show the t-SNE visualization of 2000 activated label-level visual representations ( $y_{ij} = 1$ ). In a total of 12 superclasses, the samples of most classes are discretely scattered in the original space. As shown



**Figure 4:** Visualization of the learned category prototypes on the MS-COCO dataset. Different colors and shapes mean different superclasses.



**Figure 5:** Visualization analysis of baseline method, ADDGCN method, and our method SADCL.



**Figure 6:** Visualization of the 2000 learned label-level visual representations randomly sampled images of the MS-COCO test dataset.

in Figure 6, each class representation is clustered in its own relatively independent region after using our dual contrastive learning strategy.

## 5 Conclusion

In this paper, we propose a novel semantic-aware dual contrastive learning model with semantic-aware representation learning, sample-to-sample contrastive learning, and prototype-to-sample contrastive learning. By intervening in SARL, our proposed method focuses on local regions of related categories and learns label-level visual representations. The collaborative learning of SSCL and PSCL can model complex intra-category and inter-category relationships so that the label-level feature representation with positive labels can be close to the category prototype. Experiments on five popular benchmarks demonstrate the effectiveness and state-of-the-art performance of our method.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No.61876002, 62076005, U20A20398), the Natural Science Foundation of Anhui Province (No.2008085UD07, 2008085MF191), the University Synergy Innovation Program of Anhui Province, China (No.GXXT-2021-002, GXXT-2021-030, GXXT-2021-065, GXXT-2022-029), the National Natural Science Foundation of China and Anhui Provincial Key Research and Development Project(No.202104a07020029).

## References

- [1] Tianshui Chen, Liang Lin, Riquan Chen, Xiaolu Hui, and Hefeng Wu, 'Knowledge-guided multi-label few-shot learning for general image recognition', *IEEE T-PAMI*, **44**(3), 1371–1384, (2020).
- [2] Tianshui Chen, Muxin Xu, Xiaolu Hui, Hefeng Wu, and Liang Lin, 'Learning semantic-specific graph representation for multi-label image recognition', in *ICCV*, pp. 522–531, (October 2019).
- [3] Zhao-Min Chen, Quan Cui, Borui Zhao, Renjie Song, Xiaoqin Zhang, and Osamu Yoshie, 'Sst: Spatial and semantic transformers for multi-label image recognition', *IEEE T-IP*, **31**, 2570–2583, (2022).
- [4] Zhao-Min Chen, Xiu-Shen Wei, Xin Jin, and Yanwen Guo, 'Multi-label image recognition with joint class-aware map disentangling and label correlation embedding', in *ICME*, pp. 622–627, (2019).
- [5] Zhao-Min Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo, 'Multi-label image recognition with graph convolutional networks', in *CVPR*, pp. 5177–5186, (June 2019).
- [6] Zhaomin Chen, Xiu-Shen Wei, Peng Wang, and Yanwen Guo, 'Learning graph convolutional networks for multi-label recognition and applications', *IEEE T-PAMI*, 1–1, (2021).
- [7] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng, 'Nus-wide: a real-world web image database from national university of singapore', in *ACM CIVR*, p. 48, (2009).
- [8] Son D Dao, Ethan Zhao, Dinh Phung, and Jianfei Cai, 'Multi-label image classification with contrastive learning', *arXiv*, (2021).
- [9] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman, 'The pascal visual object classes challenge: A retrospective', *IJCA*, **111**, 98–136, (2015).
- [10] Bin-Bin Gao and Hong-Yu Zhou, 'Learning to discover multi-class attentional regions for multi-label image recognition', *IEEE T-IP*, **30**, 5920–5932, (2021).
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick, 'Momentum contrast for unsupervised visual representation learning', in *CVPR*, pp. 9729–9738, (2020).
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, 'Deep residual learning for image recognition', in *CVPR*, pp. 770–778, (June 2016).
- [13] Xiao-Yuan Jing, Fei Wu, Zhiqiang Li, Ruimin Hu, and David Zhang, 'Multi-label dictionary learning for image annotation', *TIP*, **25**(6), 2712–2725, (2016).
- [14] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan, 'Supervised contrastive learning', in *Neurips*, volume 33, pp. 18661–18673, (2020).
- [15] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei, 'Visual genome: Connecting language and vision using crowdsourced dense image annotations', *IJCV*, **123**(1), 32–73, (2017).
- [16] Jack Lanchantin, Tianlu Wang, Vicente Ordonez, and Yanjun Qi, 'General multi-label image classification with transformers', in *CVPR*, pp. 16478–16488, (2021).
- [17] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi, 'Prototypical contrastive learning of unsupervised representations', in *ICLR*, (2021).
- [18] Qing Li, Xiaojiang Peng, Yu Qiao, and Qiang Peng, 'Learning label correlations for multi-label image recognition with graph networks', *PRL*, **138**, 378–384, (2020).
- [19] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick, 'Microsoft coco: Common objects in context', in *ECCV*, pp. 740–755, (2014).
- [20] Ruyang Liu, Hao Liu, Ge Li, Haodi Hou, Tinghao Yu, and Tao Yang, 'Contextual debiasing for visual recognition with causal mechanisms', in *CVPR*, pp. 12755–12765, (2022).
- [21] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu, 'Query2label: A simple transformer way to multi-label classification', *arXiv preprint arXiv:2107.10834*, (2021).
- [22] Lixiang Ru, Yibing Zhan, Baosheng Yu, and Bo Du, 'Learning affinity from attention: End-to-end weakly-supervised semantic segmentation with transformers', in *CVPR*, pp. 16846–16855, (June 2022).
- [23] Karen Simonyan and Andrew Zisserman, 'Very deep convolutional networks for large-scale image recognition', in *ICLR*, (2015).
- [24] Laurens Van der Maaten and Geoffrey Hinton, 'Visualizing data using t-sne', *JMLR*, **9**(11), (2008).
- [25] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu, 'Cnn-rnn: A unified framework for multi-label image classification', in *CVPR*, pp. 2285–2294, (June 2016).
- [26] Zhouxia Wang, Tianshui Chen, Guanbin Li, Ruijia Xu, and Liang Lin, 'Multi-label image recognition by recurrently discovering attentional regions', in *ICCV*, pp. 464–472, (Oct 2017).
- [27] Yunchao Wei, Wei Xia, Min Lin, Junshi Huang, Bingbing Ni, Jian Dong, Yao Zhao, and Shuicheng Yan, 'Hcp: A flexible cnn framework for multi-label image classification', *IEEE T-PAMI*, **38**(9), 1901–1907, (2016).
- [28] Xiangping Wu, Qingcai Chen, Wei Li, Yulun Xiao, and Baotian Hu, 'Adahgcn: Adaptive hypergraph neural networks for multi-label image classification', in *ACM MM*, pp. 284–293, (2020).
- [29] Yanan Wu, He Liu, Songhe Feng, Yi Jin, Gengyu Lyu, and Zizhang Wu, 'Gm-mlc: Graph matching based multi-label image classification', in *IJCAI*, pp. 1179–1185, (2021).
- [30] Jiazhi Xu, Sheng Huang, Fengtao Zhou, Luwen Huangfu, Daniel Zeng, and Bo Liu, 'Boosting multi-label image classification with complementary parallel self-distillation', in *IJCAI*, pp. 1495–1501, (2022).
- [31] Hao Yang, Joey Tianyi Zhou, Yu Zhang, Bin-Bin Gao, Jianxin Wu, and Jianfei Cai, 'Exploit bounding box annotations for multi-label object recognition', in *CVPR*, pp. 280–288, (June 2016).
- [32] Jie Yang, Jiarou Fan, Yiru Wang, Yige Wang, Weihao Gan, Lin Liu, and Wei Wu, 'Hierarchical feature embedding for attribute recognition', in *CVPR*, pp. 13055–13064, (June 2020).
- [33] Jin Ye, Junjun He, Xiaojiang Peng, Wenhao Wu, and Yu Qiao, 'Attention-driven dynamic graph convolutional network for multi-label image recognition', in *ECCV*, pp. 649–665, (August 2020).
- [34] Renchun You, Zhiyao Guo, Lei Cui, Xiang Long, Yingze Bao, and Shilei Wen, 'Cross-modality attention with semantic graph embedding for multi-label classification', in *AAAI*, volume 34, pp. 12709–12716, (2020).
- [35] Jiawei Zhao, Ke Yan, Yifan Zhao, Xiaowei Guo, Feiyue Huang, and Jia Li, 'Transformer-based dual relation graph for multi-label image recognition', in *CVPR*, pp. 163–172, (October 2021).
- [36] Fengtao Zhou, Sheng Huang, and Yun Xing, 'Deep semantic dictionary learning for multi-label image classification', in *AAAI*, volume 35, pp. 3572–3580, (2021).