

Multi-label image classification using adaptive graph convolutional networks: From a single domain to multiple domains

Inder Pal Singh ^{a,*}, Enjie Ghorbel ^{a,b}, Oyebade Oyedotun ^a, Djamila Aouada ^a

^a Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Luxembourg

^b Cristal Laboratory, National School of Computer Sciences, University of Manouba, Tunisia

ARTICLE INFO

Communicated by Nicu Sebe

MSC:

41A05

41A10

65D05

65D17

Keywords:

Multi-label image classification

Unsupervised domain adaptation

Machine learning

Deep learning

Computer vision

Graph Convolutional Networks

Domain shift

ABSTRACT

This paper proposes an adaptive graph-based approach for multi-label image classification. Graph-based methods have been largely exploited in the field of multi-label classification, given their ability to model label correlations. Specifically, their effectiveness has been proven not only when considering a single domain but also when taking into account multiple domains. However, the topology of the used graph is not optimal as it is pre-defined heuristically. In addition, consecutive Graph Convolutional Network (GCN) aggregations tend to destroy the feature similarity. To overcome these issues, an architecture for learning the graph connectivity in an end-to-end fashion is introduced. This is done by integrating an attention-based mechanism and a similarity-preserving strategy. The proposed framework is then extended to multiple domains using an adversarial training scheme. Numerous experiments are reported on well-known single-domain and multi-domain benchmarks. The results demonstrate that our approach achieves competitive results in terms of mean Average Precision (mAP) and model size as compared to the state-of-the-art. The code will be made publicly available.

1. Introduction

Multi-label image classification has been widely investigated by the computer vision community, given its practical relevance in numerous application fields, such as human attribute recognition (Li et al., 2016), scene classification (Shao et al., 2015) and deepfake detection (Singh et al., 2023). In contrast to traditional image classification approaches that associate a single label to an input image, multi-label image classification aims to detect the presence of a set of objects.

Thanks to the latest progress in deep learning, most of recent methods rely on a single-stream Deep Neural Networks (DNN), including Convolution Neural Networks (CNN) (Simonyan and Zisserman, 2015; He et al., 2016; Zhu et al., 2017; Ge et al., 2018; Gao and Zhou, 2021; Ridnik et al., 2021b), Residual Neural Networks (RNN) (Wang et al., 2016) and Transformers (Lanchantin et al., 2021; Cheng et al., 2022; Wang et al., 2022a; Chen et al., 2020). Nevertheless, their impressive performance comes at the cost of very large architectures that are unsuitable for memory-constrained environments. As an alternative, another line of research has tried to integrate priors related to label correlations (Chen et al., 2019b,a; Li et al., 2019; Wang et al., 2020b; Singh et al., 2022b; Wang et al., 2022b; Sun et al., 2022a). As demonstrated in Singh et al. (2022b), such an approach contributes to improving

scalability. In other words, fewer parameters are required to achieve comparable performance with traditional DNNs.

Graph-based approaches (Chen et al., 2019b; Singh et al., 2022b), are among the most popular multi-label classification methods that aim at modeling label correlations. In addition to a standard DNN that extracts image features, a second stream based on a Graph Convolutional Network (GCN) is used for generating inter-dependent label classifiers. In particular, the input graph is used to model label dependencies. Each node represents a label, and each edge is characterized by the probability of co-occurrence of a label pair. Graph-based methods have also been shown to be successful in the challenging context of cross-domain multi-label classification (Lin et al., 2021). Standard multi-label image classification methods such as Chen et al. (2019b) and Singh et al. (2022b) usually assume that unseen images and training data are drawn from the same distribution, i.e. the same domain, hence ignoring a possible domain shift problem (Lin et al., 2021). This leads, therefore, to poor generalization capabilities under cross-domain settings. UDA is a plausible solution to overcome this challenge without relying on costly annotation efforts (Ganin and Lempitsky, 2015). UDA aims at learning domain invariant features to bridge the gap between a *source* domain and a *target* domain without access to the associated labels.

* Corresponding author.

E-mail address: inder.singh@uni.lu (I.P. Singh).

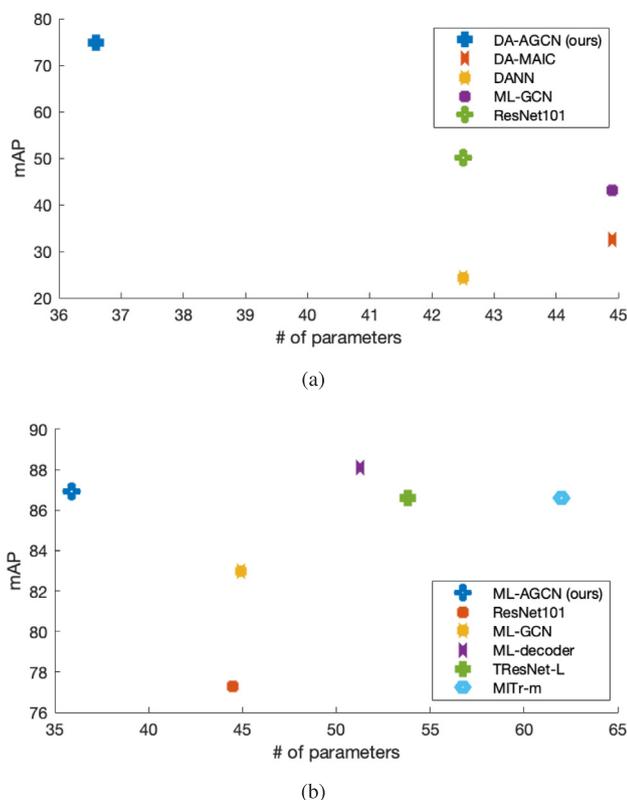


Fig. 1. Comparison of our approach (ML-AGCN) without (top) and with UDA (down) to recent state-of-the-art methods in terms of number of parameters (millions) and mean Average Precision (mAP) on MS-COCO and Clipart \rightarrow VOC. The considered state-of-the-art methods are: MITr-m (Cheng et al., 2022), TResNet-L (Ridnik et al., 2021b), ML-Decoder (Ridnik et al., 2023), ML-GCN (Chen et al., 2019b), ResNet101 (He et al., 2016), DA-MAIC (Lin et al., 2021), and DANN (Ganin et al., 2016).

In particular, one of the most successful UDA approaches for multi-label classification, namely DA-MAIC (Lin et al., 2021), leverages graph representations to model label inter-dependencies and couple it with an adversarial training approach (Ganin et al., 2016). While most existing multi-label classification methods can be extended to the cross-domain setting by just adding a simple discriminator, they mostly require a high number of parameters to work effectively. Hence, the use of a graph-based method for modeling label correlation is an interesting way for obtaining compact yet effective models, as highlighted in Lin et al. (2021). This allows achieving a good compromise between performance and compactness under the cross-domain setting.

Despite their usefulness in both single and cross-domain settings, graph-based methods (Chen et al., 2019b,a; Li et al., 2019; Wang et al., 2020b; Singh et al., 2022b; Wang et al., 2022b; Sun et al., 2022a; Lin et al., 2021) are unfortunately subject to three major limitations, namely: (1) The graph structure is heuristically defined. In particular, it is computed based on the co-occurrence of labels in the training data. Hence, this topology might not be ideal for the specific task of multi-label image classification; (2) A threshold is empirically fixed for discarding edges with a low co-occurrence probability. This means that infrequent co-occurrences are assumed to be noisy. Although this might be true in many cases, assuming that any rare event corresponds to noise does not always hold; and (3) it has been proven in Jin et al. (2021) that successive aggregation operations in the GCN usually dissipate the node similarity in the original feature space, hence potentially leading to a decrease in terms of performance.

Herein, we posit that by integrating adequate mechanisms in graph-based approaches for addressing the aforementioned issues, it should be possible to reduce the network size even more while achieving competitive performance in both single-domain and cross-domain settings.

In this paper, we propose an adaptive graph-based multi-label classification method called Multi-Label Adaptive Graph Convolutional Network (ML-AGCN) for both contexts, single domain and across domains. Our idea consists in: (1) learning two additional adjacency matrices in an end-to-end manner instead of solely relying on a heuristically defined graph topology. Note that no threshold is applied, avoiding the loss of weak yet relevant connections. In particular, the first learned graph topology computes the importance of each node pair. This is carried out by employing an attention mechanism similar to Graph Attention Networks (GAT) (Velickovic et al., 2017). Nevertheless, even though learned, the latter does not ensure the conservation of feature similarity. Hence, the second graph structure is built based on the similarity between node features and overcomes the information loss happening through successive convolutions; (2) integrating the proposed adaptive graph-based architecture in an adversarial domain adaptation framework for aligning a labeled source domain to an unlabeled target domain. As shown in Fig. 1, the results suggest that our method is competitive with respect to the state-of-the-art in terms of both mean Average Precision (mAP) and network size under the single domain and cross-domain settings.

This paper is an extended version of Singh et al. (2022a). In comparison to Singh et al. (2022a), the main contributions of this article are given below:

1. More extensive experiments showing the relevance of the proposed architecture called Multi-Label Adaptive Graph Convolutional Network (ML-AGCN) for multi-label image classification.
2. An adversarial Domain Adaptation approach integrating an Adaptive Graph Convolutional Network (DA-AGCN) for multi-label image classification.
3. A deep qualitative and quantitative experimental analysis of the proposed domain adaptation method with respect to the state-of-the-art.

The remainder of this paper is organized as follows. Section 2 reviews the state-of-the-art on multi-label image classification and domain adaptation for multi-label classification. Section 3 formulates the problem of graph-based multi-label image classification and its applicability to different domains. In Sections 4 and 5, the proposed method is detailed. Section 6 presents the experimental results and analysis. Finally, Section 7 concludes this work and highlights interesting future directions.

2. Related works

In this section, we start by presenting the state-of-the-art of multi-label image classification. Then, we focus on reviewing existing domain adaptation approaches for multi-label image classification.

2.1. Multi-label image classification

This subsection presents various works proposed for Multi-Label Image Classification (MLIC) within a single-domain context. We categorize these methodologies into two main groups based on their intuition and network architectures.

2.1.1. Single-stream deep neural networks

As discussed in Section 1, most multi-label classification methods (Razavian et al., 2014; Wei et al., 2015; Lanchantin et al., 2021; Cheng et al., 2022; Wang et al., 2022a; Ridnik et al., 2021b,a) employ a single-stream DNN. More specifically, they mainly take inspiration from successful architectures proposed in the context of single-label image classification such as CNN, RNN, and transformers. For example, Razavian et al. (2014) have used a pre-trained OVERFeat Sermanet et al. (2014) model and have adapted it to multi-label image classification. Wei et al. (2015) have leveraged the prediction of multiple CNN architectures pretrained on ImageNet (Deng et al., 2009) such

as AlexNet (Krizhevsky et al., 2012) and VGG-16 (Simonyan and Zisserman, 2015). Ridnik et al. (2021a) have employed TResNet (Ridnik et al., 2021b) using a novel loss called Asymmetric Loss (ASL) that focuses more on positive labels than negative ones. TResNet introduced in Ridnik et al. (2021b) is based on a ResNet architecture with a series of modifications for optimizing the GPU network capabilities while maintaining the performance.

Recently, Lanchantin et al. (2021) attempted to leverage transformers for modeling complex dependencies among visual features and labels. Similarly, MITr (Cheng et al., 2022) combines the pixel attention and the attention among image patches to better excavate the transformer's activity in multi-label image classification. More recently, ML-Decoder (Ridnik et al., 2023) proposed a transformer-based classification head instead of the standard Global Average Pooling (GAP), for improving the generalization capability. On the other hand, another recent work introduced a novel attention module called Interventional Dual Attention (IDA) (Liu et al., 2022) that aims to mitigate contextual bias in visual recognition through multiple sampling interventions.

2.1.2. Multi-stream deep neural networks

Going deeper into the network enables the model to learn more abstract features from the image. However, this comes at the cost of high memory requirements. Moreover, the aforementioned single-stream methods do not explicitly model the relationship between labels which can be an important semantic element to consider.

In order to incorporate the information of label correlations, a second class of methods have used a second subnetwork in addition to the main backbone. For instance, Zhu et al. (2017) have introduced a Spatial Regularization Net (SRN) to learn the underlying relationships between labels by generating label-wise attention maps. Similarly, Qu et al. (2023) have employed image representations from intermediate convolutional layers to model both local and global label semantics.

Graphs can also be an interesting way for modeling label correlations (Chen et al., 2019b; Singh et al., 2022b; Wang et al., 2020a, 2022a), while keeping the network size reasonable. ML-GCN (Chen et al., 2019b) was among the pioneering works that utilized a Graph Convolutional Network (GCN) to learn interdependent label-wise classifiers and combine them with a standard DNN that learns discriminative image features. Similarly, IML-GCN (Singh et al., 2022b) makes use of a similar GCN-CNN architecture. However, they suggest employing image-based embeddings as node features of the input graph as a replacement to the glove-based word embeddings (Pennington et al., 2014) used in Chen et al. (2019b). Indeed, GLOVE embeddings might be inappropriate for MLIC since they have been initially proposed for representing words in the field of Natural Language Processing (NLP), while images are intrinsically different.

As discussed in Section 1, despite their proven performance in terms of both precision and network size, these graph-based approaches including ML-GCN and IML-GCN suffer from some weaknesses. First, the input graph is predefined based on label co-occurrences in the training set. As a result, the heuristically defined topology labels might be sub-optimal for MLIC. Second, a threshold is set empirically in order to ignore weak edges. This induces that rare co-occurrences are presumed to be noisy. Even though this assumption often holds, some infrequent events might occur in a real-life scenario. Lastly, as discussed in Jin et al. (2021), the similarity of node features through multiple GCN layers might fade, potentially resulting in a performance decrease. The proposed ML-AGCN aims at solving these issues by adaptively learning the graph topology in an end-to-end manner, while leveraging an adequate mechanism for preserving the feature node similarity.

2.2. Unsupervised domain adaptation for multi-label image classification

Over the last years, DL methods have achieved a remarkable progress in the field of computer vision and pattern recognition. However, the effectiveness of DL approaches heavily depends on the availability of a large amount of annotated data. For mitigating the huge

cost caused by data annotation, the field of unsupervised domain adaptation (Li et al., 2020; Zhang et al., 2019; Long et al., 2017; Ganin et al., 2016; Long et al., 2018) has been widely investigated over the last decade. It aims at making use of an existing labeled dataset from a related domain called *source domain* to enhance the model performance on a domain of interest termed *target domain*, for which only unlabeled data are provided.

Unsupervised domain adaptation methods can be separated into two main categories. The first one (Li et al., 2020; Zhang et al., 2019; Long et al., 2017) aims at explicitly reducing the domain gap by minimizing statistical discrepancy measures between the two domains. Alternatively, the second class of methods implicitly minimizes this domain gap by adopting an adversarial training approach (Ganin et al., 2016; Long et al., 2018). The main idea consists in using a domain classifier to play a min-max two-player game with the feature generator. This strategy is designed to enforce the generation of domain-invariant features that are sufficiently discriminative.

Nevertheless, most existing techniques focus on the task of single-label image classification. In fact, very few papers have considered domain adaptation for multi-label image classification (Li et al., 2021; Pham et al., 2021; Chen et al., 2019b; Singh et al., 2024).

Among these rare references, we can mention ML-ANet (Li et al., 2021), which explicitly minimizes the domain gap by optimizing multi-kernels maximum mean discrepancies (MK-MMD) in a Reproducing Kernel Hilbert Space (RKHS). More recently, an adversarial approach has been adopted in Pham et al. (2021) where a condition-based domain discriminator similar to conditional-GANs (Mirza and Osindero, 2014) has been employed. However, similar to the first category of methods for traditional multi-label image classification, these two approaches (Li et al., 2021; Pham et al., 2021) neglect the important information of label dependencies.

A graph-based approach called DA-MAIC has been then proposed as an alternative (Lin et al., 2021). As in ML-GCN (Chen et al., 2019b), they have proposed to build a graph for modeling label correlations based on label co-occurrences. Additionally, to reduce the domain shift between the source and target domains, they have used a domain classifier that is trained in an adversarial manner. Unfortunately, DA-MAIC is impacted by the same drawbacks affecting ML-GCN (Chen et al., 2019b) as detailed in Sections 1 and 2.1. More details regarding these issues are given in the next section.

Inspired by Chen et al. (2022), a very recent work (Singh et al., 2024) called DDA-MLIC redefining the adversarial loss based on the multi-label classifier has been proposed, hence eliminating the need for an additional discriminator. Although this approach has shown promising results, it addresses the problem of domain adaptation from a different perspective. Indeed, the present paper aims to propose a suitable multi-label classification mechanism that could be beneficial for both single-domain and cross-domain settings while DDA-MLIC investigates a more adequate adversarial strategy. Therefore, the two methods (DDA-MLIC and ours) tackle two different yet complementary aspects of UDA for multi-label classification.

3. Background and problem formulation

In this section, Graph Convolutional Networks (GCN) are first reviewed. Then, the problem of multi-label image classification in both single and cross-domain settings is formulated.

3.1. Background: Graph Convolutional Networks (GCN)

CNNs are defined on a regular grid and are, therefore, not directly applicable to non-Euclidean structures such as graphs. GCN (Kipf and Welling, 2017) have been proposed as the generalization of traditional CNN to graphs. They have been very successful in numerous computer vision applications, including human pose estimation (Cai et al., 2019) and human action recognition (Papadopoulos et al., 2020a,b). Let us

denote a graph by $G = (V, E, \mathbf{F})$. The set $V = \{v_1, v_2, \dots, v_N\}$ is formed by N nodes, while $E = \{e_1, e_2, \dots, e_M\}$ refers to the set of M edges connecting the nodes. Finally, $\mathbf{F} = \{\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_N\}$ represents the node features such that $\mathbf{f}_i \in \mathbb{R}^d$ corresponds to the features of the node v_i .

Let us assume that \mathbf{F}^l is the input node features of the l th layer and $\mathbf{A} \in \mathbb{R}^{N \times N}$ the adjacency matrix of the graph G . Each GCN layer can be seen as a non-linear function $h(\cdot)$ that computes the node features $\mathbf{F}^{l+1} \in \mathbb{R}^{N \times d'}$ of the $(l+1)$ th layer as follows,

$$\mathbf{F}^{l+1} = h(\mathbf{A}\mathbf{F}^l\mathbf{W}^l), \quad (1)$$

where $\mathbf{W}^l \in \mathbb{R}^{d \times d'}$ is the weight matrix. Note that \mathbf{A} is normalized before using Eq. (1).

3.2. Problem formulation

3.2.1. Graph-based multi-label image classification

The goal of multi-label image classification is to predict the presence or absence of a set of objects $\mathcal{O} = \{1, 2, \dots, N\}$ in a given image \mathbf{I} . This can be done by learning a function f such that,

$$f: \mathbb{R}^{w \times h} \rightarrow \llbracket 0, 1 \rrbracket^N \quad (2)$$

$$\mathbf{I} \mapsto \mathbf{y} = [y_i]_{i \in \mathcal{O}},$$

where w and h define the pixel-wise width and height of the image, respectively, and $y_i = 1$ indicates the presence of the label i in \mathbf{I} , in contrast to $y_i = 0$. Graph-based multi-label image classification methods such as ML-GCN (Chen et al., 2019b) and IML-GCN (Singh et al., 2022b) usually involve two subnetworks: (1) A feature generator denoted by f_g , and (2) an estimator of N inter-dependent binary classifiers denoted by f_c . The generator mostly corresponds to an out-of-the-shelf CNN network which produces a d_f -dimensional image feature representation as described below,

$$f_g: \mathbb{R}^{w \times h} \rightarrow \mathbb{R}^{d_f} \quad (3)$$

$$\mathbf{I} \mapsto \mathbf{X}$$

For instance, ML-GCN integrates a ResNet101 (He et al., 2016), while IML-GCN makes use of TResNet-M (Ridnik et al., 2021b).

On the other hand, the second subnetwork f_c is a GCN formed by L layers which takes a fixed graph $G = (V, E, \mathbf{F})$ as input where $\text{card}(V) = N$. In fact, each node $v_i \in V$ refers to the label $i \in \{1, 2, \dots, N\}$ and each $\mathbf{f}_i \in \mathbf{F}$ corresponds to its associated label embedding. The adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ of G is usually pre-computed based on the co-occurrence probabilities that are estimated over the training set (Singh et al., 2022b; Chen et al., 2019b). In addition, rare co-occurrences are discarded as they are assumed to be noisy. In other words, given an empirically fixed threshold τ ,

$$\mathbf{A}_{ij} = \begin{cases} 0, & \text{if } p_{ij} < \tau, \\ 1, & \text{if } p_{ij} \geq \tau \end{cases}, \quad (4)$$

with $p_{ij} = p(I_j | I_i)$ being the probability of co-occurrence of label j and label i in the same image.

Given $\mathbf{F}^l \in \mathbb{R}^{d_l \times L}$ as the input node features of the l th layer, the input features \mathbf{F}^{l+1} of the $(l+1)$ th are therefore computed by following Eq. (1).

Finally, the vertex features produced by the last GCN layer, i.e., $\mathbf{F}^L \in \mathbb{R}^{N \times d_f}$, form the N inter-dependent classifiers.

In summary, f_c can be defined as follows,

$$f_c: \mathcal{G}(N) \rightarrow \mathbb{R}^{d_f \times N} \quad (5)$$

$$G \mapsto \mathbf{F}^L = [F_1^L, \dots, F_N^L],$$

where $\mathcal{G}(N)$ represents the set of graphs with N nodes, and \mathbf{F}_i^L for $i \in \{1, \dots, N\}$ is the generated inter-dependent binary classifier associated to the label i .

Hence, f , which returns the final prediction, can be defined as follows,

$$f(\mathbf{I}) = \text{sig}(f_c(G)f_g(\mathbf{I})^T) \quad (6)$$

$$= \text{sig}(\mathbf{F}^L \mathbf{X}^T),$$

where $\text{sig}(x) = \frac{1}{1+e^{-x}}$ is the sigmoid activation function.

However, as explained in Section 1, three main limitations can be noted in graph-based methods: (1) the computation of the adjacency matrix \mathbf{A} is made heuristically and is decoupled from the training process; (2) a threshold τ is empirically fixed for completely ignoring rare co-occurrences; and (3) as shown in Jin et al. (2021), aggregating successively node features in a graph may induce the loss of the similarity/dissimilarity information present in the initial feature space.

3.2.2. Graph-based unsupervised domain adaptation for multi-label image classification

Let us consider a source dataset for multi-label image classification $\{(\mathbf{I}_s^k, \mathbf{y}_s^k)\}_{k=1}^{n_s}$ drawn from a distribution D_s and a similar unlabeled target dataset $\{(\mathbf{I}_t^k)\}_{k=1}^{n_t}$ drawn from a different distribution D_t . The matrix $\mathbf{I}_s^k \in \mathbb{R}^{w \times h}$ refers to the k th image sample of D_s and $\mathbf{y}_s^k \in \llbracket 0, 1 \rrbracket^N$ is its associated label vector, while $\mathbf{I}_t^k \in \mathbb{R}^{w \times h}$ denotes the k th image sample of D_t . The variables n_s and n_t refer respectively to the total number of samples in D_s and D_t . The goal of unsupervised domain adaptation is to train a model using labeled source domain data sampled from D_s and unlabeled target domain samples from D_t for making accurate predictions on the target domain. Despite its relevance in real-world applications, few domain adaptation methods have been proposed for multi-label image classification. Among the most recent and accurate approaches, one can mention DA-MAIC (Lin et al., 2021), which also takes advantage of the graph representation for modeling label correlations. It directly extends ML-GCN (Chen et al., 2019b) by integrating an adversarial training for domain adaptation. Consequently, DA-MAIC is subject to the same limitations induced by ML-GCN (Chen et al., 2019b) mentioned in Section 3.2.1. Our assumption is that by finding the appropriate solutions to the enumerated issues, we can, at the same time, enhance the performance of existing multi-label image classification and domain adaptation for multi-label image classification methods.

4. Multi-label Adaptive Graph Convolutional Network (ML-AGCN)

To handle the challenges mentioned in Section 3, a novel graph-based approach called Multi-Label Adaptive Graph Convolutional Network (ML-AGCN) is introduced.

4.1. Overview of the proposed architecture

Similar to Chen et al. (2019b) and Singh et al. (2022b), a network formed by two subnets is adopted as illustrated in Fig. 2. The first is a CNN that extracts a discriminative representation from a given input image, while the second is a GCN-based network that learns N interdependent classifiers. As in Singh et al. (2022b), TResNet-M which represents a small version of TResNet (Ridnik et al., 2021b) is employed as a CNN subnetwork. TResNet is a direct extension of ResNet, which fully exploits the GPU capabilities to boost the model efficiency. The graph-based subnet, Adaptive Graph Convolutional Network (AGCN), uses the same image embeddings proposed in Singh et al. (2022b) as feature nodes. Nevertheless, in contrast to Singh et al. (2022b) and Chen et al. (2019b), it relies on an end-to-end learned graph topology. More details regarding this subnetwork are provided in Section 4.2. Similar to Singh et al. (2022b), the Asymmetric Loss (ASL) (Ridnik et al., 2021a) denoted by \mathcal{L}_c is used for optimizing ML-AGCN such that,

$$\mathcal{L}_c = \mathbb{E}_{(\mathbf{I}_s, \mathbf{y}_s) \sim D_s} \sum_{i=1}^N y_s^{(i)} (1 - p^{(i)})^{\gamma_+} \log(p^{(i)}) + (1 - y_s^{(i)}) (p_m^{(i)})^{\gamma_-} \log(1 - p_m^{(i)}), \quad (7)$$

where γ_+ and γ_- are focusing parameters for positive and negative samples, respectively, and $y_s^{(i)}$ and $p^{(i)}$ are the respective ground truth and predicted probability with respect to the label i and $p_m^{(i)}$ is the shifted probability given by $\max(p^{(i)} - m, 0)$, where m is a threshold used for reducing the effect of easy negative samples (Ridnik et al., 2021a).

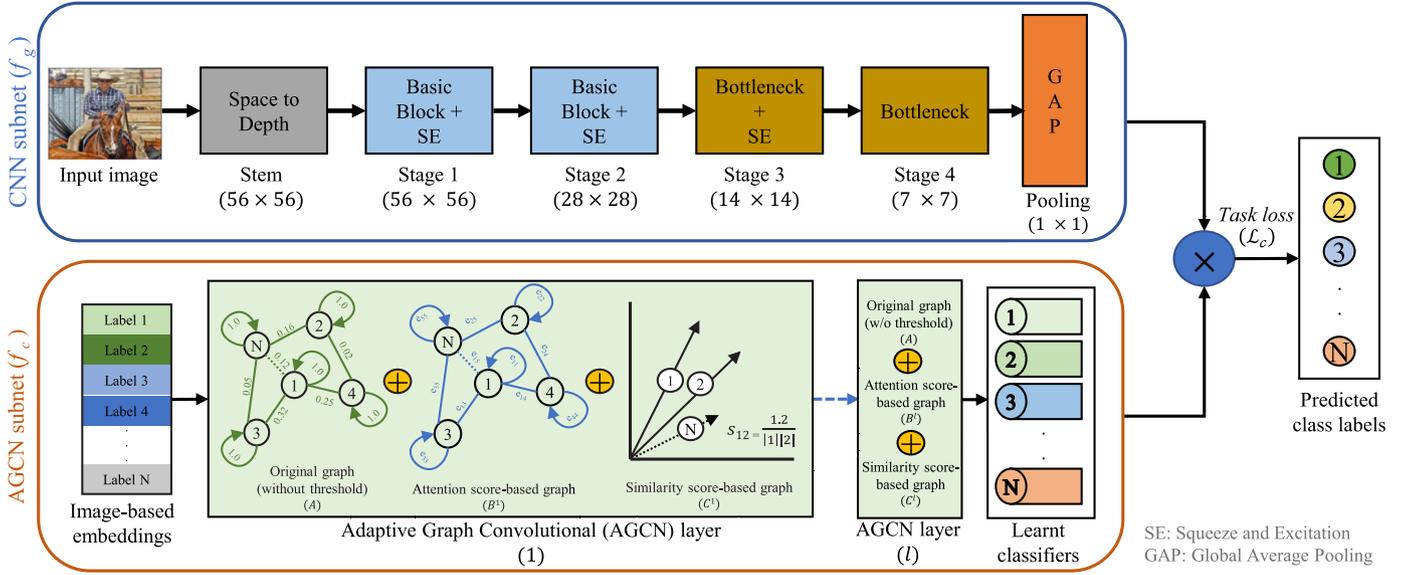


Fig. 2. Architecture of ML-AGCN (Singh et al., 2022a): On the one hand, the CNN subnet learns relevant image features from an input image. On the other hand, the GCN subnet estimates interdependent label classifiers by taking into account one fixed adjacency matrix A and two adaptive adjacency matrices $B^{(l)}$ and $C^{(l)}$. Finally, the classifiers are applied to the CNN features for predicting the labels.

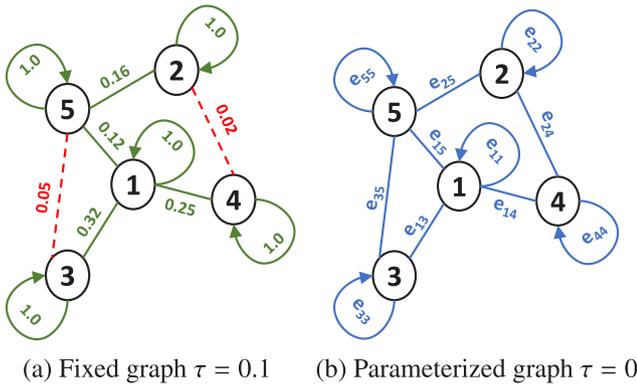


Fig. 3. (a) An example of a fixed label graph with a threshold set to $\tau = 0.1$ (Singh et al., 2022a). Dashed (red) edges indicate the ignored edges; (b) The proposed parameterized graph topology considering all the edges. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.2. Graph-based subnet: Adaptive Graph Convolutional Network (AGCN)

Our intuition is that by integrating a suitable mechanism, it should be possible to boost the classification performance and reduce at the same time the size of graph-based methods. Hence, as illustrated in Fig. 3, we propose to adaptively learn the graph topology by reformulating Eq. (1) as below,

$$\mathbf{F}^{l+1} = \sigma((\mathbf{A} + \mathbf{B}^{(l)} + \mathbf{C}^{(l)})\mathbf{F}^l\mathbf{W}^l), \quad (8)$$

where σ is a LeakyRELU activation function.

Instead of relying solely on the adjacency matrix A defined in Chen et al. (2019b), two additional parameterized graphs called attention-based and similarity adjacency graphs, respectively denoted by $B^{(l)}$ and $C^{(l)}$, are defined. In this case, no threshold is applied to A for ignoring rare co-occurrences. It is also important to note that A is fixed, while $B^{(l)}$ and $C^{(l)}$ vary from one layer to another. In the following, we detail how $B^{(l)}$ and $C^{(l)}$ are computed.

4.2.1. Attention-based adjacency matrix

Instead of ignoring rare co-occurrences in A , the matrix $B^{(l)} = (b_{ij}^{(l)})_{i,j \in \mathcal{O}}$ is defined based on an attention mechanism where the importance of each edge is quantified. To that aim, inspired by Velickovic et al. (2017), an attention score denoted by e_{ij} is calculated for each pair of vertices (v_i, v_j) as follows,

$$e_{ij} = \sigma(\mathbf{a}^{(l)T} (\mathbf{W}\mathbf{F}_i^{(l)} \parallel \mathbf{W}\mathbf{F}_j^{(l)})), \quad (9)$$

where $\mathbf{W} \in \mathbb{R}^{d^{(l+1)} \times d^{(l)}}$ represents a learnable weight matrix, $\mathbf{a}^{(l)T} \in \mathbb{R}^{2d^{(l+1)}}$ are the learnable attention coefficients and \parallel refers to the concatenation operation.

A softmax function is then applied to the computed normalized attention scores such that,

$$\alpha_{ij}^{(l)} = \frac{\exp(e_{ij}^{(l)})}{\sum_{k \in \mathcal{N}(i)} \exp(e_{ik}^{(l)})}, \quad (10)$$

with $\mathcal{N}(i)$ defining the neighborhood of the node i and α_{ij} being the obtained normalized attention score.

We recall that the goal of the GCN subnet is to generate interdependent label classifiers. This means that each classifier must be predominated by the information related to the label it belongs to. Hence, the attention score of the node in question should be maximal. For this purpose, an additional step called self-importance mechanism is proposed for computing the attention-based adjacency matrix $B^{(l)} = (b_{ij}^{(l)})_{i,j \in \mathcal{O}}$ as follows,

$$\begin{cases} b_{ij}^{(l)} = \alpha_{ij}^{(l)} + \max_{k \in \mathcal{O}}(\alpha_{ik}^{(l)}) & \text{if } i = j \\ b_{ij}^{(l)} = \alpha_{ij}^{(l)} & \text{if } i \neq j. \end{cases} \quad (11)$$

4.2.2. Similarity-based adjacency matrix (C)

As illustrated in the top row of Fig. 4, the GCN aggregation process tends to modify the node similarity in the original feature space (Jin et al., 2021). Subsequently, the second row illustrates a significant change in the similarity measure between the node features before and after the GCN operation (see Fig. 4(a) and (b), respectively). Although learned in an end-to-end manner, the attention-based graph B only quantifies the connectivity importance of each node pair and does not guarantee that the feature similarity is not lost through graph

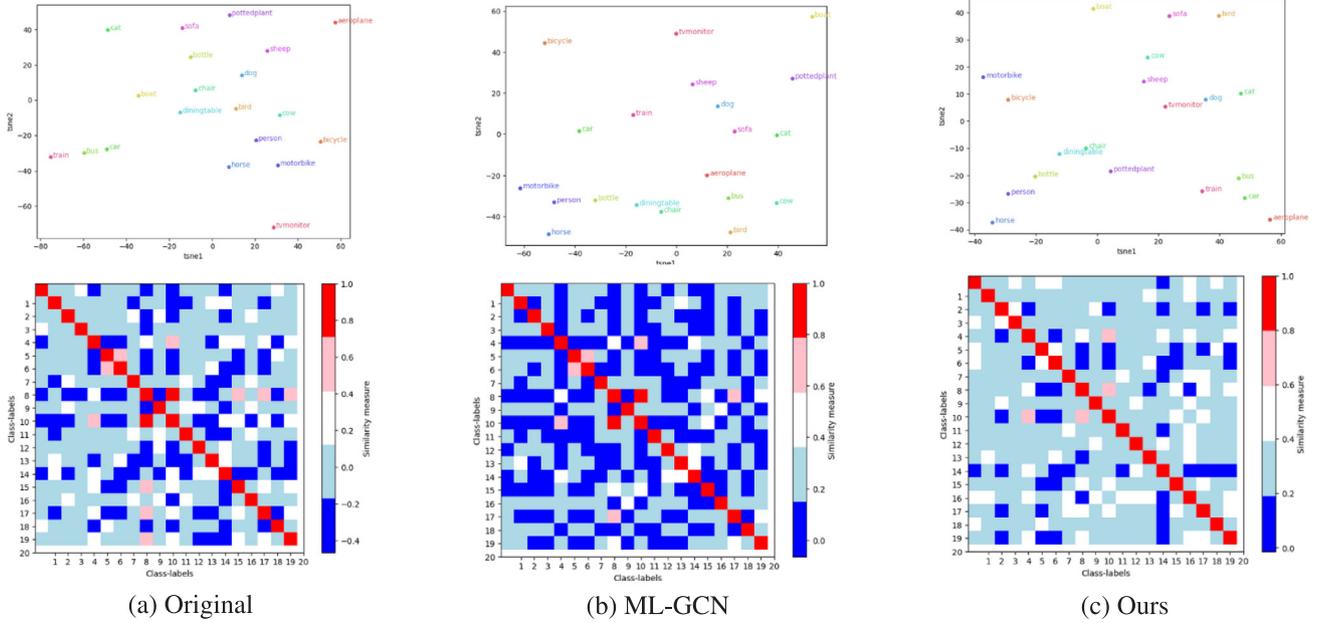


Fig. 4. Comparison of node feature similarity: The top row presents a tSNE visualization, while the bottom row illustrates cosine-similarity map between the graph nodes for VOC dataset: (a) using the original image-based embeddings (before GCN), (b) after applying two layers of standard GCN using the proposed architecture in ML-GCN (Chen et al., 2019b), and (c) after applying two layers of AGCN using our approach (i.e., ML-AGCN).

convolutions. Hence, we propose to incorporate an additional matrix $\mathbf{C}^{(l)} = (c_{ij}^{(l)})_{i,j \in \mathcal{O}}$ to preserve the node feature similarity. It is obtained by calculating the cosine similarity $c_{ij}^{(l)}$ between each pair of vertices (v_i, v_j) as follows,

$$c_{ij}^{(l)} = \frac{\mathbf{F}_i^{(l)} \cdot \mathbf{F}_j^{(l)}}{\|\mathbf{F}_i^{(l)}\| \|\mathbf{F}_j^{(l)}\|}, \quad (12)$$

where $\|\cdot\|$ denotes the L_2 Euclidean norm.

Finally, the output of the final layer L denoted by \mathbf{F}^L is used in Eq. (6) for predicting the labels. As shown in Fig. 4(c), by using this strategy, the node similarity is relatively preserved, as compared to a standard GCN.

5. Unsupervised domain adaptation for multi-label image classification using ML-AGCN

The proposed architecture for multi-label image classification called DA-AGCN is illustrated in Fig. 5. Similar to ML-AGCN, we make use of TRResNet-based backbone f_g to extract discriminative image features and an Adaptive Graph Convolutional Network f_c to learn N inter-dependent classifiers. Given source and target input images from D_s and D_t , respectively, the goal of f_g is to generate D -dim domain-invariant image features. The latter is used to predict the image labels, as depicted in Eq. (6). The label classification loss \mathcal{L}_c is therefore defined as in Eq. (7).

For obtaining domain-invariant representations, a domain classifier $f_d: \mathbb{R}^D \rightarrow \{0, 1\}$ is employed and trained in an adversarial manner. Given the features $\mathbf{X} = f_g(\mathbf{I})$ extracted from a given image \mathbf{I} , f_d predicts the domain of the input image as follows,

$$f_d: \mathbb{R}^D \rightarrow \{0, 1\} \\ \mathbf{X} \mapsto \hat{d}. \quad (13)$$

where \hat{d} is the predicted domain label of \mathbf{I} . Note that the ground-truth domain label $d = 0$ if \mathbf{I} is from the source domain and $d = 1$ if sampled from the target domain.

As in Ganin et al. (2016), a domain loss is defined as below,

$$\mathcal{L}_d = \mathbb{E}_{f_g(\mathbf{I}_s) \sim D_s} \log \frac{1}{f_d(f_g(\mathbf{I}_s))} +$$

$$\mathbb{E}_{f_g(\mathbf{I}_t) \sim D_t} \log \frac{1}{(1 - f_d(f_g(\mathbf{I}_t)))}. \quad (14)$$

Given \mathcal{L}_c defined in Eq. (7), the final objective function used to optimize the network is $E(\theta_g, \theta_d, \theta_c)$ defined by,

$$E(\theta_g, \theta_d, \theta_c) = \mathcal{L}_c(\theta_c, \theta_g) + \lambda \mathcal{L}_d(\theta_d, \theta_g), \quad (15)$$

where θ_g, θ_d and θ_c refer, respectively, to the parameters of f_g, f_d and f_c and λ is a hyper-parameter defining the weight of \mathcal{L}_d . The network is trained in an adversarial manner using the GRL for obtaining the optimal parameters $(\hat{\theta}_g, \hat{\theta}_c, \hat{\theta}_d)$ such that,

$$(\hat{\theta}_g, \hat{\theta}_c) = \min_{\theta_g, \theta_c} E(\theta_g, \theta_c, \hat{\theta}_d), \\ \hat{\theta}_d = \max_{\theta_d} E(\hat{\theta}_g, \hat{\theta}_c, \theta_d). \quad (16)$$

6. Experiments

In this section, the experimental settings and results are presented and discussed for: (1) single-domain multi-label image classification (Section 6.1); and (2) domain-adaptation for multi-label image classification (Section 6.2).

6.1. Multi-label image classification

6.1.1. Implementation details

As discussed in Section 4, TRResNet-M (Ridnik et al., 2021b) is used as the CNN backbone. In particular, the fully connected layer after the Global Average Pooling (GAP) layer is removed. The output of the GAP layer produces a 2048-dimensional latent image representation. The dimension of the AGCN output has also been set to 2048. For the node features, we use the image-based embeddings proposed in Singh et al. (2022b) instead of the usual word-based embeddings used in Chen et al. (2019b). The model is trained for 40 epochs using Adam, with a maximum learning rate of $1e-4$ using a cosine decay.

6.1.2. Datasets

In the following, the datasets that have been used for the experiments are presented.

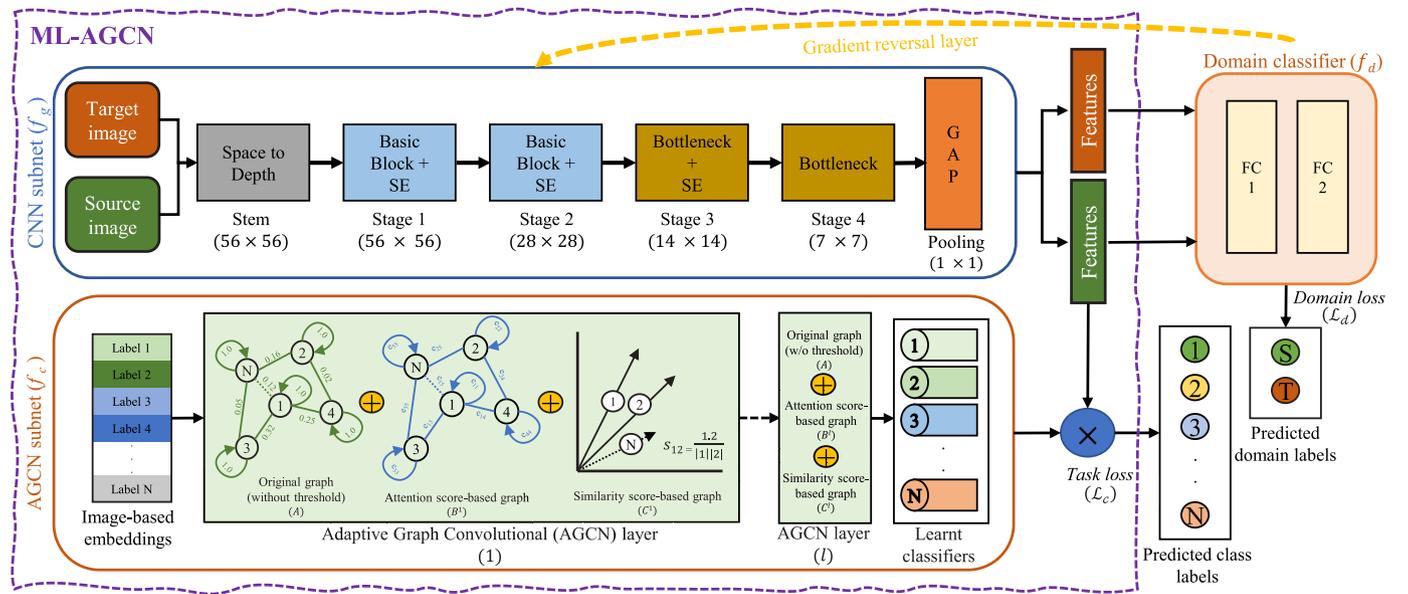


Fig. 5. Architecture of the proposed DA-AGCN for multi-label image classification (best viewed in color). Images from both source and target datasets are given as input to the CNN subnet that generates image features. The AGCN-subnet, similar to ML-AGCN (Singh et al., 2022a), learns in an end-to-end manner the attention and similarity-based adjacency matrices B^0 and C^0 , respectively, and generates accordingly inter-dependent label classifiers using only labeled source images. In addition, a domain classifier is considered. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

MS-COCO. MS-COCO (Lin et al., 2014) is a widely used large-scale multi-label image dataset. It contains 80K training images and 40k testing images. Each image is annotated with multiple object labels from a total of 80 categories.

VG-500. The VG-500 dataset (Chen et al., 2020) is a well-known dataset for multi-label image classification. It includes 500 different objects as categories. The dataset comes with a training set of 98,249 images and a testing set of 10,000 images.

PASCAL-VOC 2007. The PASCAL Visual Object Classes Challenge (Everingham et al., 2010) introduced in 2007 is one of the most commonly used multi-label image classification datasets. It contains about 10K image samples with 5011 and 4952 images as training and testing sets, respectively. The images show 20 different object categories with an average of 2.5 categories per image.

6.1.3. Quantitative analysis

Comparison with state-of-the-art methods in terms of mAP and model size.

We compare the performance of ML-AGCN with current state-of-the-art methods by reporting the mean Average Precision (mAP) as well as the number of model parameters. Additionally, similar to Singh et al. (2022b) and Chen et al. (2019b), we report the following evaluation metrics on the MS-COCO dataset, namely, average per-Class Precision (CP), average per-Class Recall (CR), average per-Class F1-score (CF1), average Overall Precision (OP), average overall recall (OR) and average Overall F1-score (OF1).

Tables 1–3 report the quantitative comparison of the proposed approach with respect to state-of-the-art methods on the MS-COCO, the VG-500, and the VOC datasets, respectively. It can be clearly seen that our method achieves competitive results as compared to existing methods in terms of mAP while considerably reducing the model size. More specifically, similar to SSGRL, we achieve the best mAP performance on VOC-2007 while reducing the number of parameters from 92.2 to 35.8 million. Moreover, ML-AGCN reaches the second-best mAP performance on MS-COCO and VG-500. However, as indicated in Tables 1 and 2, our method achieves the second-best performance on MS-COCO and VG-500. Specifically, ML-decoder (Ridnik et al., 2023) and C-Tran (Lanchantin et al., 2021) outperform our approach by 1.2% and 0.5% in mAP on MS-COCO and VG-500. This slight difference in

performance might be explained by the fact that ML-decoder and C-Tran incorporate a noticeably higher number of parameters, i.e., 51.3 and 120 million, respectively, against only 35.9 and 32.1 million for ML-AGCN. Consequently, our approach achieves comparable performance while necessitating around 30% and 70% less parameters than ML-decoder and C-Tran, respectively. Such results show that our approach can maintain competitive performance while considerably reducing the network as compared to the state-of-the-art.

It is also worth mentioning that ML-AGCN outperforms the ML-GCN (Chen et al., 2019b) baseline by 3.9%, 6.3%, and 0.5% in terms of mAP on the MS-COCO, VG and VOC datasets, respectively, while keeping a comparable number of parameters.

Moreover, for a fair comparison with this baseline, we re-train ML-GCN by replacing the original CNN backbone (ResNet101) with the same architecture used in our experiments, namely, TRResNetM. The results that are reported in Tables 1, 2, and 3 show that ML-AGCN outperforms ML-GCN with a higher margin when using the same backbone, thereby confirming the relevance of the proposed adaptive graph convolution module.

In addition, we also report that the obtained performance when including only one layer in the graph-subnet of ML-GCN (Chen et al., 2019b), IML-GCN (Singh et al., 2022b) and ML-AGCN. The obtained results confirm the relevance of the proposed adaptive learning. In fact, it can be seen that our method outperforms existing graph-based methods under this setting. More specifically, ML-AGCN achieves an improvement of 5.7% and 5.3% in terms of mAP when compared to ML-GCN and IML-GCN, respectively, on the MS-COCO dataset. Similarly, on the VG-500 benchmark, a significant increase of 20.7% is recorded in terms of mAP in comparison to IML-GCN.

Ablation study. In order to analyze the impact of each adjacency matrix used in our approach, namely, the attention-based matrix B and similarity-based matrix C , an ablation study is carried out as shown in Table 4. It can be noted that by considering B in addition to A (without threshold), an improvement of 5.1%, 20.5%, and 0.04% in terms of mAP is made, respectively, on MS-COCO, VG-500, and VOC-2007. The magnitude of improvement seems dependent on the number of classes contained in the considered dataset. In fact, while VG-500 is formed by 500 classes, MS-COCO and VOC-2007 are respectively composed of 80 and 20 categories. This highlights the importance of

Table 1

Comparison with the state-of-the-art methods on the MS-COCO dataset. (Best and second-best performances are indicated in bold and are underlined, respectively).

Category	Method	Resolution	# params	mAP	CP	CR	CF1	OP	OR	OF1
CNN	Multi-Evidence (Ge et al., 2018)	448 × 448	49.0	–	80.4	70.2	74.9	85.2	72.5	78.4
	SRN (Zhu et al., 2017)	224 × 224	48.0	77.1	81.6	65.4	71.2	82.7	69.9	75.8
	ResNet101 (He et al., 2016)	448 × 448	44.5	77.3	80.2	66.7	72.8	83.9	70.8	76.8
	MCAR (Gao and Zhou, 2021)	448 × 448	44.9	83.8	85	72.1	78	88	73.9	80.3
	MCAR (Gao and Zhou, 2021)	576 × 576	44.9	84.5	84.3	73.9	78.7	86.9	76.1	81.1
	TResNet-L (Ridnik et al., 2021a)	448 × 448	53.8	86.6	<u>87.4</u>	76.4	<u>81.4</u>	<u>88.1</u>	79.2	81.8
RNN	CNN-RNN (Wang et al., 2016)	224 × 224	66.2	61.2	–	–	–	–	–	–
Graph-based	ML-GCN (1-layer) (Chen et al., 2019b)	448 × 448	43.1	80.9	82.9	69.7	75.8	84.8	73.6	78.8
	IML-GCN (1-layer) (Singh et al., 2022b)	448 × 448	29.5	81.3	81.3	72.2	76.0	86.7	77.9	82.1
	ML-GCN (Chen et al., 2019b)	448 × 448	44.9	83.0	85.1	72.0	78.0	85.8	75.4	80.3
	ML-GCN (TResNetM) (Chen et al., 2019b)	448 × 448	31.9	82.4	<u>87.6</u>	66.6	75.2	<u>91.0</u>	70.3	79.3
	A-GCN (Li et al., 2019)	512 × 512	44.1	83.1	84.7	72.3	78.0	85.6	75.5	80.3
	F-GCN (Wang et al., 2020b)	–	44.3	83.2	85.4	72.4	78.3	86.0	75.7	80.5
	CFMIC (Wang et al., 2022b)	448 × 448	45.1	83.8	85.8	72.7	78.7	86.3	76.3	81.0
	SSGRL (Chen et al., 2019a)	576 × 576	92.2	83.8	89.9	68.5	76.8	91.3	70.8	79.7
	FLNet (Sun et al., 2022a)	–	46.0	84.1	84.9	73.9	79.0	85.5	77.4	81.1
	IML-GCN (Singh et al., 2022b)	448 × 448	31.5	86.6	78.8	82.6	80.2	79.0	85.1	81.9
Transformer-based	IDA (Liu et al., 2022)	576 × 576	55.1	86.3	–	–	80.4	–	–	82.5
	MITr-s (Cheng et al., 2022)	384 × 384	33.0	83.9	82.8	75.5	77.3	83	78.5	79.9
	STMG (Wang et al., 2022a)	384 × 384	197.0	84.3	85.8	72.7	78.7	86.7	76.8	81.5
	C-Tran (Lanchantin et al., 2021)	576 × 576	120.0	85.1	86.3	74.3	79.9	87.7	76.5	81.7
	MITr-m (Cheng et al., 2022)	384 × 384	62.0	86.8	84	80.1	81.7	84.6	82.5	<u>83.5</u>
	ML-Decoder (Ridnik et al., 2023)	448 × 448	51.3	88.1	–	–	–	–	–	–
	Ours: ML-AGCN (1-layer)	448 × 448	<u>29.9</u>	86.7	79.6	<u>82.4</u>	80.7	79.8	<u>84.5</u>	82.1
Ours: ML-AGCN (2-layers)	448 × 448	35.9	<u>86.9</u>	86.2	78.3	81.7	87.2	80.7	83.8	

Table 2

Comparison with the state-of-the-art methods on the VG-500 dataset. (Best and second to the best performance is indicated in bold and underline respectively).

Method	# params	mAP	CP	CR	CF1	OP	OR	OF1
ML-GCN (TResNetM) (Chen et al., 2019b)	33.2	26.6	30.7	7.7	10.3	75.4	10.0	17.7
ResNet101 (He et al., 2016)	44.5	30.9	39.1	25.6	31.0	61.4	35.9	45.4
ML-GCN (Chen et al., 2019b)	44.9	32.6	42.8	20.2	27.5	66.9	31.5	42.8
TResNet-M (Ridnik et al., 2021a)	29.5	33.6	–	–	–	–	–	–
IML-GCN (Singh et al., 2022b)	<u>32.1</u>	34.5	–	–	–	–	–	–
SSGRL (Chen et al., 2019a)	92.2	36.6	–	–	–	–	–	–
KGGR (Chen et al., 2020)	45.0	37.4	<u>47.4</u>	24.7	32.5	66.9	36.5	47.2
C-Tran (Lanchantin et al., 2021)	120.0	38.4	49.8	<u>27.2</u>	35.2	66.9	<u>39.2</u>	<u>49.5</u>
Ours: ML-AGCN (1-layer)	<u>32.1</u>	<u>37.9</u>	47.2	31.8	<u>34.7</u>	64.1	42.1	50.8
Ours: ML-AGCN (2-layers)	37.4	37.1	47.3	24.7	29.0	<u>67.6</u>	38.1	48.7

Table 3

Comparison with the state-of-the-art methods on the VOC-2007 dataset. (Best and second to the best performance is indicated in bold and underline respectively).

Methods	# params	mAP	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Table	Dog	Horse	mbike	Person	Plant	Sheep	Sofa	Train	Tv
CNN-RNN (Wang et al., 2016)	66.2	84.0	96.7	83.1	94.2	92.8	61.2	82.1	89.1	94.2	64.2	83.6	70	92.4	91.7	84.2	93.7	59.8	93.2	75.3	99.7	78.6
VeryDeep (Simonyan and Zisserman, 2015)	138	89.7	98.9	95.0	96.8	95.4	69.7	90.4	93.5	96.0	74.2	86.6	87.8	96.0	96.3	93.1	97.2	70.0	92.1	80.3	98.1	87.0
ResNet101 (He et al., 2016)	44.5	89.9	99.5	97.7	97.8	96.4	65.7	91.8	96.1	97.6	74.2	80.9	85	98.4	96.5	95.9	98.4	70.1	88.3	80.2	98.9	89.2
HCP (Wei et al., 2015)	138.0	90.9	98.6	97.1	98.0	95.6	75.3	94.7	95.8	97.3	73.1	90.2	80.0	97.3	96.1	94.9	96.3	78.3	94.7	76.2	97.9	91.5
ML-GCN (Chen et al., 2019b)	44.9	94.0	99.5	<u>98.5</u>	98.6	98.1	80.8	94.6	97.2	98.2	82.3	95.7	86.4	98.2	98.4	96.7	<u>99.0</u>	84.7	96.7	84.3	98.9	93.7
ML-GCN (TResNetM) (Chen et al., 2019b)	31.8	94.1	<u>99.8</u>	98.0	98.9	98.0	80.4	95.9	96.0	97.8	83.4	<u>98.6</u>	86.4	98.6	99.0	95.1	98.8	82.7	98.8	84.5	99.7	92.0
F-GCN (Wang et al., 2020b)	–	94.1	99.5	<u>98.5</u>	98.7	<u>98.2</u>	80.9	94.8	97.3	98.3	82.5	95.7	86.6	98.2	98.4	96.7	<u>99.0</u>	84.8	96.7	84.4	99.0	93.7
FLNet (Sun et al., 2022a)	46.0	94.4	99.6	98.1	98.9	97.9	84.6	95.3	96.2	96.5	85.6	96.1	87.2	97.7	98.6	<u>97.0</u>	98.1	<u>86.5</u>	97.4	86.5	98.8	90.8
TResNet-L (Ridnik et al., 2021a)	53.8	94.6	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
CFMIC (Wang et al., 2022b)	45.1	94.7	99.7	<u>98.5</u>	<u>98.8</u>	98.3	83.9	96.5	<u>97.5</u>	<u>98.8</u>	83.1	96.1	87.4	98.6	<u>98.9</u>	97.2	<u>99.0</u>	85.4	97.1	84.9	<u>99.2</u>	94.2
MCAR (Gao and Zhou, 2021)	44.9	<u>94.8</u>	99.7	99.0	98.5	<u>98.2</u>	<u>85.4</u>	96.9	97.4	98.9	<u>83.7</u>	95.5	<u>88.8</u>	<u>99.1</u>	98.2	95.1	99.1	84.8	97.1	<u>87.8</u>	98.3	94.8
SSGRL (Chen et al., 2019a)	92.2	95.0	99.7	98.4	98.0	97.6	85.7	96.2	98.2	98.8	82.0	98.1	89.7	98.8	98.7	<u>97.0</u>	<u>99.0</u>	86.9	<u>98.1</u>	85.8	99.0	93.7
Ours ML-AGCN (1-layer)	29.5	94.5	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–
Ours ML-AGCN (2-layers)	35.8	95.0	99.9	98.0	98.5	98.0	81.6	<u>96.8</u>	96.6	98.2	85.6	99.4	88.2	99.2	99.0	96.5	98.8	84.8	99.5	88.1	98.9	<u>94.5</u>

Table 4

The ablation of adaptively learning **B** and **C** for multi-label image classification in single domain.

Method	Mean Average Precision (mAP %)		
	MS-COCO	VG-500	VOC-2007
ML-AGCN (A)	81.1	17	94.27
ML-AGCN (A+B)	86.6 (+5.1%)	37.5 (+20.5%)	94.31 (+0.04%)
ML-AGCN (A+B+C)	86.7 (0.1%)	37.9 (+0.4%)	95.0 (+0.69%)

Table 5

Hyper-parameter analysis: comparison of the performance of ML-AGCN when varying the number of AGCN layers. Best performance is indicated in **bold**.

# layers	# params (millions)	mean Average Precision (%)		
		MS-COCO	VG-500	VOC-2007
1-layer	29.9	86.7	37.9	94.5
2-layers	35.9	86.9	37.1	95.0
3-layers	37.3	85.7	24.6	94.7
4-layers	40.7	85.4	–	94.6

adaptively modeling label correlations, especially when dealing with a large number of classes, which is likely to occur in a practical scenario. An additional mAP improvement of 0.1%, 0.4%, and 0.17% can be seen when including **C** in the AGCN subnet. In conclusion, it can be noted that **B** contributes more importantly to the resulting enhancement. This could be explained by the fact that **C** is less needed, as only two layers are considered in the graph subnet. As shown in Fig. 4., while the use of **C** allows to better preserve the feature similarity as, the latter is not completely lost through layers when using 2 standard GCN convolutions.

Hyper-parameter analysis. Table 5 reports the mean Average Precision (mAP) for MS-COCO, VG-500, and VOC-2007 datasets, when varying the total number of layers in ML-AGCN. The best results are generally obtained when only two layers are considered. However, it should be noted that the results are very stable for MS-COCO and VOC-2007, with a variation lower than 1.5% for all the configurations.

6.1.4. Qualitative analysis

In Fig. 6, the Gradient-weighted Class Activation Mapping (Grad-CAM) (Selvaraju et al., 2017) is visualized for some examples from the VOC dataset. While the first column of images represents the input images, the second, third, and fourth columns show, respectively, the Grad-cam visualization from a model trained using only **A**, **A** and **B**, and finally **A**, **B** and **C**. These qualitative results are in line with the quantitative ones. As shown in Fig. 6, the use of **B** allows activating more precisely the regions of interest in the image, while the contribution of **C** in this refinement is less impressive but remains visible.

6.2. Unsupervised domain adaptation for multi-label image classification

6.2.1. Implementation details

We reproduce the results of current state-of-the-art methods due to the limited availability of DA approaches for multi-label image classification. In particular, we first consider standard multi-label image classification methods (without DA) and refer to them as *MLIC*. Since no target images were employed during the learning process, this is equivalent to source-only training. Second, we evaluate the zero-shot performance of existing Vision Language Models (VLM) on the target dataset referred to as VLM in our experiments. For this purpose, we use a pre-trained Contrastive Language-Image Pre-Training CLIP (Radford et al., 2021) model with a ViT-B/32 transformer-based backbone. As commonly done in MLIC (Sun et al., 2022b), we compute a similarity score, given an input image to VLM, between the predicted text features and ground truth features of the considered labels. The application of

a sigmoid activation function to these scores results in the presence probabilities of the considered objects.

Lastly, we reproduce the outcomes of two state-of-the-art domain adaptation methods, namely DANN (Ganin et al., 2016) and DA-MAIC (Lin et al., 2021), that we refer as *DA* in our experiments. Moreover, we replace the multi-label softmargin-loss with the traditional cross-entropy loss when training DANN. We generate Glove-based word embeddings (Pennington et al., 2014) as node features for training DA-MAIC. Additionally, in order to showcase the effectiveness of the proposed AGCN subnet, we present the results of DANN and DA-MAIC by considering the same backbone as ours: an additional experiment based on TResNet-M instead of the traditional ResNet101 is carried out. The images have been resized to 224×224 , unless stated differently. Our domain classifier includes one hidden layer of dimension 1024. A maximum learning rate of $1e-4$ using a cosine decay is considered. The model is trained for a total of 40 epochs or until convergence.

6.2.2. Datasets

Similar to DA-MAIC (Lin et al., 2021), we use three multi-label aerial image datasets in our experiments, namely, AID (Hua et al., 2020), UCM (Chaudhuri et al., 2018) and DFC15 (Hua et al., 2019). Additionally, due to the limited number of suitable datasets for the task of DA in MLIC, we convert two well-known object detection datasets, initially used for DA in the context of object detection, to multi-label annotations, namely PASCAL VOC 2007 (Everingham et al., 2010) and Clipart1k (Inoue et al., 2018).

AID multi-label aerial dataset. The original AID dataset (Xia et al., 2017) contains 10 000 high-resolution aerial images. The images cover a total of 30 categories. A multi-label version of this dataset was produced in Hua et al. (2020), where 3000 aerial images from the original AID dataset have been selected and assigned with multiple object labels. In total, they include 17 labels: airplane, sand, pavement, buildings, cars, chaparral, court, trees, dock, tank, water, grass, mobile-home, ship, bare-soil, sea, and field. 80% and 20% of the images have been used respectively for training and testing (Hua et al., 2020).

UCM multi-label aerial dataset. UCM multi-label dataset (Chaudhuri et al., 2018) is derived from the UCM dataset (Yang and Newsam, 2010). It consists of images showing 21 land-use classes. The images have a resolution of 256×256 pixels. Later in Chaudhuri et al. (2018), 2100 of these aerial images were annotated with multiple tags in order to generate a multi-label aerial image dataset. This dataset shares the same number of labels as AID multi-label dataset (Hua et al., 2020) i.e., 17 labels. In our experiments, 80% and 20% of data are respectively used for training and testing.

DFC15 multi-label aerial dataset. The DFC15 multi-label dataset (Hua et al., 2019) was initially introduced in 2015. It has a total of 3342 high-resolution image samples and includes 8 object labels. In our experiments, the 6 categories in common with UCM and AID datasets, including water, grass, building, tree, ship, and car, are considered. 80% and 20% are respectively used for training and testing.

VOC and Clipart1k datasets. The Clipart1k (Inoue et al., 2018) dataset contains 20 object categories, similar to PASCAL-VOC 2007 (Everingham et al., 2010). We create a multi-label annotation for each image by considering the category of each object bounding box. Clipart1k provides a total of 1000 image samples. 50% and 50% are respectively used for training and testing.

6.2.3. Quantitative analysis

Comparison with state-of-the-art methods in terms of mAP and model size. The proposed domain adaptation approach for multi-label image classification is compared with state-of-the-art methods. The same metrics described in Section 6.1.3 are used, including mAP, CP, CR, CFI, OP, OR, and OF1. In addition to the four protocols followed in DA-MAIC (Lin et al., 2021), i.e., AID \rightarrow UCM, UCM \rightarrow AID, AID \rightarrow DFC,

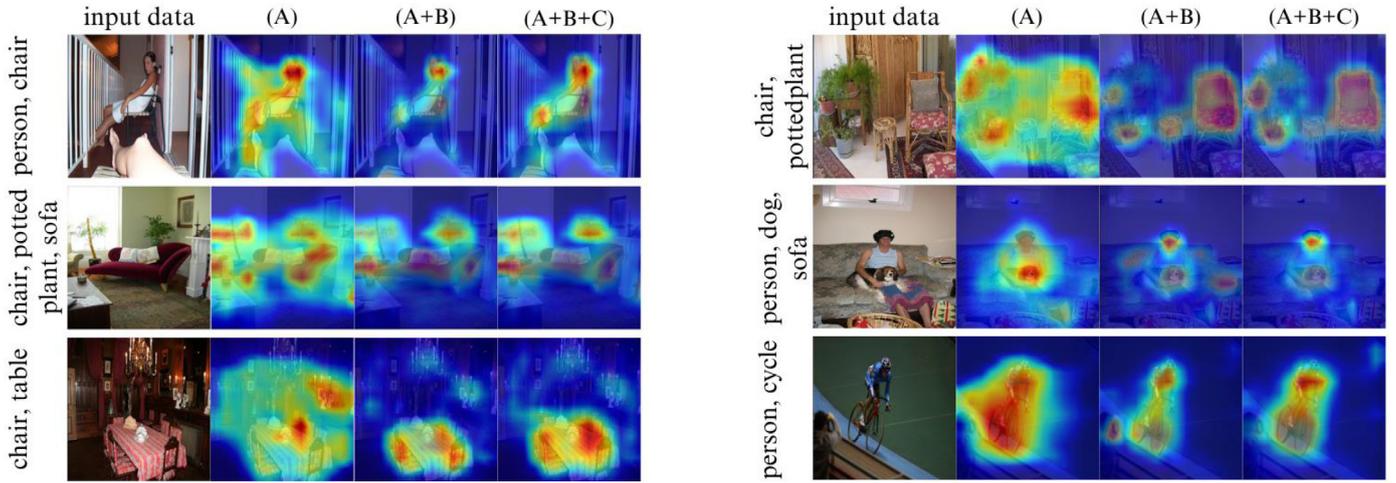


Fig. 6. Grad-CAM visualization of the predictions with ML-AGCN (Singh et al., 2022a) using samples from the VOC dataset using only A, then A and B, and finally A, B and C.

Table 6

Comparison with the state-of-the-art in terms of mAP and number of model parameters using two settings, i.e., AID \rightarrow UCM and UCM \rightarrow AID. (Best performance is indicated in bold and second-best performance is underlined).

Category	Backbone	Method	# params	AID \rightarrow UCM							UCM \rightarrow AID						
				mAP	CP	CR	CF1	OP	OR	OF1	mAP	CP	CR	CF1	OP	OR	OF1
MLIC	ResNet50	RESNET (He et al., 2016)	23.5	54.5	57.9	42.3	43.2	<u>70.2</u>	67.1	68.6	50.3	57.2	26.8	32.2	<u>90.0</u>	44.2	59.3
	ResNet101	RESNET (He et al., 2016)	42.5	57.5	60.0	47.5	47.0	69.1	71.5	70.3	51.7	50.6	29.6	33.9	88.0	48.5	62.5
		ML-GCN (Chen et al., 2019b)	44.9	53.7	55.3	44.3	45.9	70.2	68.7	<u>69.4</u>	51.3	50.1	29.9	34.0	88.0	49.7	63.6
	TResNet-M	ML-GCN (Chen et al., 2019b)	31.8	53.5	57.8	40.5	41.1	64.2	70.4	67.2	52.5	51.1	25.8	31.9	91.6	40.8	56.5
		ASL (Ridnik et al., 2021a)	29.4	55.4	48.7	52.8	47.1	58.7	79.1	67.4	54.1	54.5	40.2	<u>41.9</u>	85.4	65.1	73.9
	ML-AGCN (Singh et al., 2022a)	36.6	55.2	36.6	64.9	45.1	45.0	88.1	59.6	52.1	48.2	47.4	42.9	77.1	79.8	78.4	
VLM	ViT-B/32	CLIP (Radford et al., 2021)	151.3	42.2	45.8	25.2	28.2	51.2	23.6	32.3	42.4	44.8	21.8	22.4	54.1	15.9	24.5
DA	ResNet101	DANN (Ganin et al., 2016)	42.5	55.3	54.7	57.9	<u>52.8</u>	59.0	81.6	68.5	50.5	<u>60.4</u>	24.9	32.3	89.2	42.2	57.3
		DA-MAIC (Lin et al., 2021)	44.9	49.7	52.4	48.7	45.4	56.9	72.5	63.8	48.7	51.6	<u>40.9</u>	41.5	78.9	65.5	71.6
	TResNet-M	DANN (Ganin et al., 2016)	29.4	52.5	<u>59.1</u>	31.6	36.3	70.9	53.7	61.1	51.6	52.1	23.2	27.9	83.2	27.8	41.7
		DA-MAIC (Lin et al., 2021)	31.8	54.4	55.3	37.5	38.6	68.0	67.9	67.9	50.5	51.8	22.9	29.0	91.6	35.2	50.8
		DDA-MLIC (Singh et al., 2024)	29.4	63.2	52.5	<u>63.7</u>	55.1	59.4	<u>82.8</u>	69.2	<u>54.9</u>	53.9	30.4	35.5	84.6	41.0	55.3
	DA-AGCN (Ours)	36.6	<u>59.0</u>	54.0	59.4	52.3	57.1	82.3	67.4	57.2	61.2	40.0	41.8	86.1	<u>66.9</u>	<u>75.3</u>	

Table 7

Comparison with the state-of-the-art in terms of mAP and number of model parameters using two settings, i.e., UCM \rightarrow DFC and AID \rightarrow DFC. (Best performance is indicated in bold and second-best performance is underlined).

Category	Backbone	Method	# params	AID \rightarrow DFC							UCM \rightarrow DFC						
				mAP	CP	CR	CF1	OP	OR	OF1	mAP	CP	CR	CF1	OP	OR	OF1
MLIC	ResNet50	RESNET (He et al., 2016)	23.5	52.9	52.3	47.1	44.2	52.3	50.5	51.3	67.7	58.5	31.9	34.9	65.9	39.1	49.1
	ResNet101	RESNET (He et al., 2016)	42.5	56.9	52.9	61.5	48.7	46.1	63.7	53.5	66.4	74.4	31.2	36.9	67.2	37.2	47.9
		ML-GCN (Chen et al., 2019b)	44.9	58.9	56.7	57.9	45.8	45.7	65.0	53.7	64.6	72.4	32.0	35.6	64.4	38.9	48.5
	TResNet-M	ML-GCN (Chen et al., 2019b)	31.8	53.5	52.0	47.9	41.6	<u>52.7</u>	51.1	51.9	66.6	60.2	35.0	38.1	64.9	39.4	49.0
		ASL (Ridnik et al., 2021a)	29.4	56.1	49.6	68.4	49.9	43.5	74.1	54.8	68.9	66.3	53.1	44.0	52.6	57.0	54.7
	ML-AGCN (Singh et al., 2022a)	36.6	51.6	41.5	83.8	52.3	40.2	88.7	55.3	70.3	68.4	<u>56.1</u>	47.8	53.8	<u>58.5</u>	56.0	
VLM	ViT-B/32	CLIP (Radford et al., 2021)	151.3	<u>64.3</u>	88.1	22.3	31.4	78.7	22.0	34.4	63.8	88.4	22.7	31.3	79.5	22.5	35.1
DA	ResNet101	DANN (Ganin et al., 2016)	42.5	<u>64.3</u>	<u>57.0</u>	65.9	51.4	48.0	65.9	55.6	67.2	72.8	44.7	48.7	56.3	49.1	52.4
		DA-MAIC (Lin et al., 2021)	44.9	50.1	<u>54.5</u>	45.8	38.9	44.2	48.8	46.4	65.6	69.7	54.2	<u>52.8</u>	57.1	58.4	<u>57.7</u>
	TResNet-M	DANN (Ganin et al., 2016)	29.4	43.0	40.7	13.6	19.3	46.0	15.6	23.3	64.1	<u>77.3</u>	22.6	30.1	<u>68.6</u>	26.5	38.2
		DA-MAIC (Lin et al., 2021)	31.8	55.4	49.8	60.4	44.7	47.3	64.1	54.4	65.8	71.4	39.3	39.7	59.9	44.6	51.1
		DDA-MLIC (Singh et al., 2024)	29.4	62.1	47.6	75.5	<u>55.3</u>	48.9	76.2	59.6	<u>70.6</u>	67.2	55.7	49.3	55.0	58.4	56.6
	DA-AGCN (Ours)	36.6	65.7	51.8	78.1	55.7	45.2	<u>80.8</u>	<u>58.0</u>	76.5	68.5	61.7	59.0	60.0	60.2	60.1	

and UCM \rightarrow DFC, two more combinations VOC \rightarrow Clipart and Clipart \rightarrow VOC are provided. Two categories of methods are considered in our evaluation: (1) the conventional Multi-Label Image Classification (MLIC); and the Domain Adaptation-based (DA) methods that aim at explicitly reducing the gap between source and target datasets.

Table 6 reports the results obtained considering AID and UCM datasets. Two different settings are followed: AID \rightarrow UCM, where AID

is the source dataset and UCM is the target dataset, and UCM \rightarrow AID, where UCM is the source dataset and AID is the target dataset. It can be clearly seen, in both types of DA settings, that the proposed DA-AGCN outperforms existing state-of-the-art for UCM \rightarrow AID. However, we are the second best in terms of mAP for AID \rightarrow UCM, where the recent discriminator-free DDA-MLIC (Singh et al., 2024) achieves a higher mAP. On the other hand, an improvement of mAP by 10.65% and

Table 8

Comparison with the state-of-the-art in terms of mAP and number of model parameters using the two settings, i.e., VOC \rightarrow Clipart and Clipart \rightarrow VOC. (Best performance is indicated in bold and second-best performance is underlined).

Category	Backbone	Method	# params	VOC \rightarrow Clipart							Clipart \rightarrow VOC						
				mAP	CP	CR	CF1	OP	OR	OF1	mAP	CP	CR	CF1	OP	OR	OF1
MLIC	ResNet50	RESNET (He et al., 2016)	23.5	42.0	57.6	15.3	22.5	82.3	25.8	39.3	45.4	40.3	9.8	13.0	84.7	25.5	39.2
	ResNet101	RESNET (He et al., 2016)	42.5	38.0	64.8	14.3	22.5	82.3	18.3	29.9	50.1	66.2	17.5	25.5	83.9	29.6	43.7
		ML-GCN (Chen et al., 2019b)	44.9	43.5	62.5	20.3	28.4	86.6	27.8	42.1	43.1	57.9	21.0	26.8	73.5	30.6	43.2
	TResNet-M	ML-GCN (Chen et al., 2019b)	31.8	53.0	75.3	26.9	37.3	90.3	31.3	46.4	67.6	80.8	39.9	50.8	88.8	46.0	60.6
		ASL (Ridnik et al., 2021a)	29.4	56.8	72.0	38.5	47.6	82.8	45.7	58.9	64.2	69.0	30.7	37.3	80.0	45.7	58.2
		ML-AGCN (Singh et al., 2022a)	36.6	53.7	75.5	35.5	44.4	79.1	39.9	53.1	38.0	45.5	25.1	28.2	61.8	36.6	45.9
VLM	ViT-B/32	CLIP (Radford et al., 2021)	151.3	58.0	43.6	57.8	43.6	45.0	42.3	43.6	73.1	53.3	76.6	60.1	53.5	<u>58.8</u>	56.0
DA	ResNet101	DANN (Ganin et al., 2016)	42.5	33.9	47.7	17.0	20.6	57.3	24.1	33.9	24.3	28.3	16.4	14.2	39.8	24.7	30.5
		DA-MAIC (Lin et al., 2021)	44.9	25.8	28.0	3.1	5.1	<u>92.6</u>	2.9	5.6	32.6	48.2	12.2	14.4	50.2	31.9	39.0
	TResNet-M	DANN (Ganin et al., 2016)	29.4	40.0	<u>82.4</u>	17.2	27.4	93.8	17.5	29.5	67.0	76.8	23.3	32.6	93.1	20.4	33.4
		DA-MAIC (Lin et al., 2021)	31.8	<u>62.3</u>	77.4	42.6	<u>51.6</u>	83.1	51.0	63.2	74.3	<u>84.5</u>	53.9	<u>63.0</u>	83.7	57.7	68.3
		DDA-MLIC (Singh et al., 2024)	29.4	61.4	84.7	28.1	39.4	90.9	33.3	48.8	77.0	86.9	29.3	38.2	88.4	35.3	50.4
	DA-AGCN (Ours)	36.6	62.9	75.8	<u>46.8</u>	54.0	80.0	<u>50.4</u>	<u>61.8</u>	<u>75.8</u>	70.9	<u>71.8</u>	69.8	66.1	75.6	70.5	

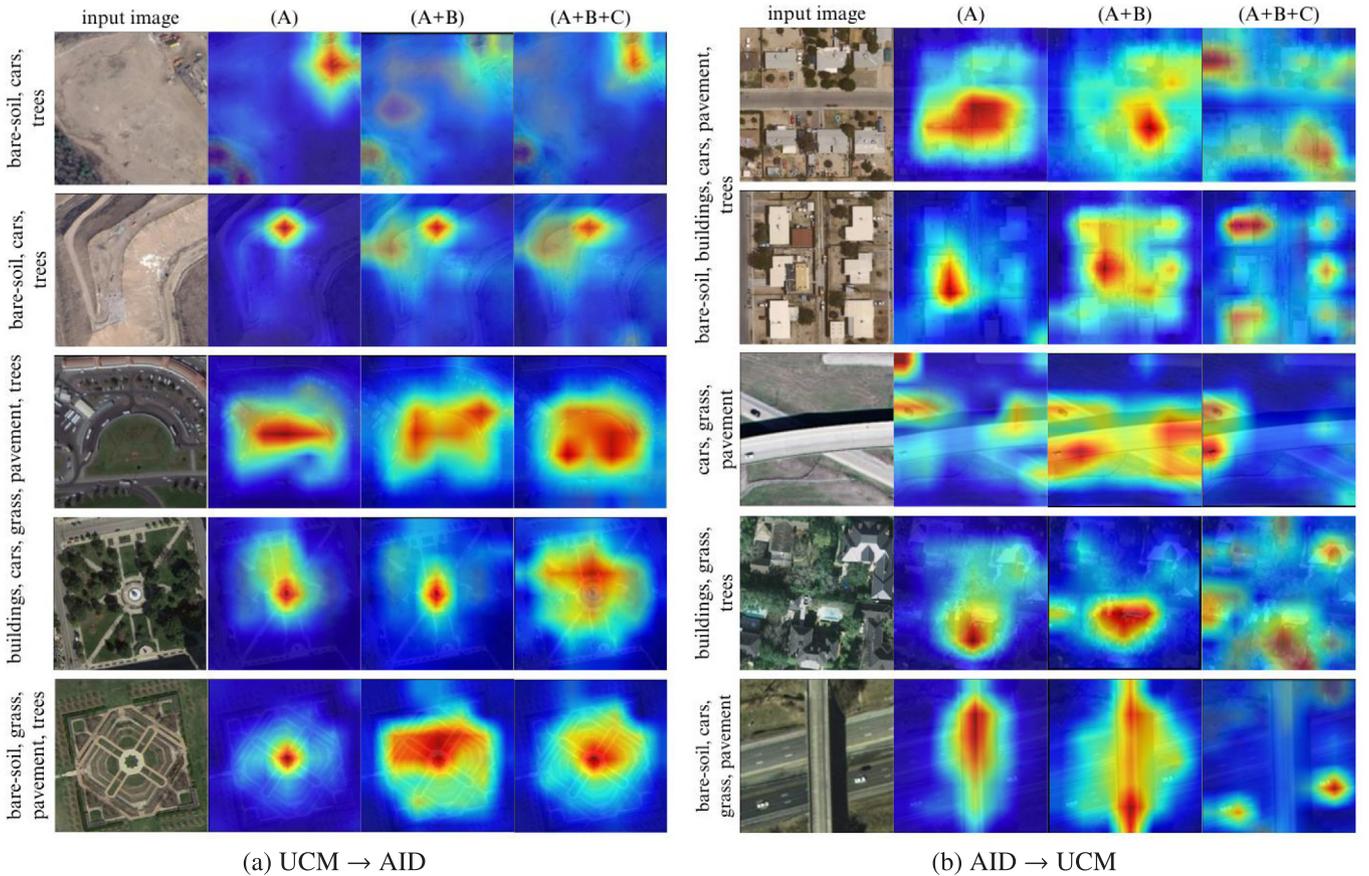


Fig. 7. Grad-CAM visualization of the multi-label predictions using the proposed DA-AGCN (a) UCM \rightarrow AID, and (b) AID \rightarrow UCM setting using only A, then A and B, and finally A, B and C.

4.59% has been respectively recorded for AID \rightarrow UCM and UCM \rightarrow AID as compared with the DA-MAIC (Lin et al., 2021) baseline. This suggests that learning the graph topology in an end-to-end manner helps improve the performance in the presence of a domain shift. Moreover, the proposed method outperforms existing state-of-the-art VLM method by more than 17% in terms of mAP despite having approximately five times the lesser number of model parameters.

In Table 7, DA-AGCN is also compared with existing methods by considering the settings UCM \rightarrow DFC and AID \rightarrow DFC. It can be clearly noticed that DA-AGCN outperforms the existing works, including the VLM and very recent DDA-MLIC in terms of mAP and model size for both UCM \rightarrow DFC and AID \rightarrow DFC. More precisely, an improvement

of 1.4% and 6.2% is achieved in terms of mAP compared to the second-best performing methods.

Table 8 also confirms the superiority of the proposed approach as compared to other state-of-the-art techniques. However, it is to note that the improvement given by the adaptive graph remains limited for Clipart \rightarrow VOC, with a slightly lower performance than DDA-MLIC. This might be explained by the fact that DA for MLIC datasets include a relatively low number of classes (8 to 20 categories). Hence, the interest of adaptively modeling label correlations might be not entirely visible. This demonstrates the necessity of creating benchmarks for DA under an MLIC context, including a wider number of classes.

Table 9
Ablation study: effect of adaptively learning **B** and **C** in the presence of a domain shift.

Graphs			mean Average Precision (mAP %)					
A	B	C	AID → UCM	UCM → AID	AID → DFC	UCM → DFC	VOC → Clipart	Clipart → VOC
✓			33.6	42.8	42.6	58.3	53.1	68.3
	✓		55.3	50.6	53.7	69.7	61.3	80.0
		✓	47.0	51.9	56.7	70.0	50.9	43.4
✓	✓		45.1	55.4	51.5	67.9	52.5	75.3
	✓	✓	54.7	52.7	60.0	71.8	58.6	76.1
✓		✓	48.1	54.5	53.7	71.6	62.0	75.6
✓	✓	✓	59.0	57.2	65.7	76.5	62.9	75.8

Table 10
Ablation study: impact of using an adversarial strategy and learning an adaptive graph topology.

Methods/Dataset (mAP)	UCM → AID	AID → UCM	AID → DFC	UCM → DFC
TResNet only	34.98	37.03	62.36	54.77
TResNet + AGCN	53.62 (+18.64)	53.08 (+16.05)	58.44 (-3.92)	74.4 (+19.68)
TResNet + AGCN + DC	55.79 (+2.17)	55.37 (+2.29)	65.11 (+6.67)	76.49 (+2.04)

In summary, our approach achieves the best performance for 4 settings over 2 and the second-best performance for the two remaining ones, ranking just after DDA-MLIC. This highlights the relevance of the proposed adaptive method under the unsupervised domain adaptation settings. In future works, it will be interesting to study the complementarity of the two best-performing approaches, namely, ours and DDA-MLIC.

Ablation study. Table 10 reports the obtained results when considering each module comprised in DA-AGCN. More specifically, the first row reports the mAP score on the target dataset using a model trained on the source dataset without incorporating the learning of label correlations and without any domain adaptation. The results in the second row are obtained by incorporating the proposed adaptive graph learning strategy for modeling the label correlations (ML-AGCN). It can be clearly seen that adaptively learning the label graph topology based on the proposed graphs **B** and **C** leads to a significant performance improvement. More specifically, an mAP improvement of at least 16 to 20% for UCM → AID, AID → UCM and UCM → DFC can be observed. Note that the results in the second row are obtained without any domain adaptation. Finally, the last row reports the results when adding an additional domain classifier to implicitly minimize the domain gap. This further improves the mAP score by an additional 2 to 6%, compared to the ML-AGCN, across all benchmarks, by this remains marginal as compared to the enhancement induced by the graph structure adaptive learning. This might return to the fact that the domain gap in aerial datasets remains relatively small, hence benefiting more from the adaptive graph strategy than the alignment.

Additionally, in Table 9, we report the performance obtained on five datasets by incorporating different combinations of **A**, **B** and **C**. It can be noted in the first row that the original adjacency matrix **A** alone cannot effectively model the label relationship, hence leading to the lowest mAP score across almost all benchmarks. The next two rows of the table further showcase that using either **B** or **C** provides a significant performance improvement when compared to only **A**, across all datasets except VOC and Clipart where **C** alone does not yield superior performance as compared to **A**. This not surprising since **C** aims to preserve the node feature similarity, and is not expected to model the label correlations. Furthermore, the next three rows show the variation in performance when coupling two to three matrices. For instance, combining either **B** or **C** with the original adjacency matrix **A** always provides a higher mAP when compared to only **A**. Finally, the last row confirms the superiority of the proposed combination **A**, **B** and **C** as the highest mAP score across all benchmarks is reached except for Clipart → VOC. A possible explanation to this exception is the synthetic nature of the Clipart dataset, where the label co-occurrence is well-thought and might be representative enough.

Furthermore, the importance of preserving the node feature similarity is demonstrated in the last row, where a notable improvement in

mAP is observed as compared to the **A** + **B** setting. As compared to the single-domain setup, the improvement resulting from **C** is more significant (see Table 4, where the contribution of **C** is relatively marginal). This might be explained by the fact that in a cross-domain setting, an adversarial training is adopted for enforcing the generation of domain-invariant features. This min-max game known for its instability, might amplify the node feature dissimilarity phenomenon.

Hyper-parameter analysis. In Table 11, we report the mean Average Precision (mAP) on the six considered benchmarks when varying the hyper-parameter (λ), introduced in Eq (15). Specifically, we vary its value from 0.1 to 1.0. It can be noted that the obtained results are relatively stable with a variation in mAP of more or less 5%.

6.2.4. Qualitative analysis

Grad-CAM (Selvaraju et al., 2017) visualization in the presence of a domain shift can be seen in Fig. 7. While Fig. 7(a) demonstrate these results for UCM → AID, Fig. 7(b) showcases the visualizations for AID → UCM. The left-most column is the input image, and the next three consecutive columns represent the Grad-CAM visualization of the image using a model trained with; **A**, **A** and **B**, and **A**, **B**, **C** respectively. It can be clearly seen that adaptively learning the graph topology **B** and **C** helps activate the most relevant areas of interest, leading to better classification performance.

6.2.5. Failure cases

In Fig. 8, some failure cases are presented using the Grad-CAM visualization. The first column shows the input image with ground truth labels while the subsequent three columns are the activated Grad-CAMs with **A**, **A** and **B** and **A**, **B** and **C**, respectively. It can be noticed that when the targeted object occupies most of the image, the use of **B** and **C** makes the model confuse the object with the background. A good illustration of this can be seen in the second row of Fig. 8. The *car* is confused with the *grass*. In future work, modeling the object occupancy will be investigated for handling these failure cases.

7. Conclusion

Existing graph-based methods have shown great performances for multi-label image classification in the context of both single-domain and cross-domain. However, these methods mostly fix the graph topology heuristically while discarding edges with rare co-occurrences. Furthermore, it has been demonstrated in Jin et al. (2021) that successive GCN aggregations tend to destroy the initial feature similarity. Hence, as a solution, an adaptive strategy for learning the graph in an end-to-end manner is proposed. In particular, attention-based and similarity-preserving mechanisms are adopted. The proposed framework for multi-label classification in a single domain is then extended to

Table 11
Sensitivity analysis: performance of the proposed DA-AGCN using different values of the hyper-parameter (λ), defined in Eq. (15).

lambda (λ)	mean Average Precision (mAP %)					
	AID→UCM	UCM→AID	AID→DFC	UCM→DFC	VOC→Clipart	Clipart→VOC
0.1	52.2	52.5	56.5	68.4	51.2	73.9
0.2	53.3	54.9	53.1	70.6	52.8	75.8
0.3	53.5	52.0	58.2	72.8	52.6	71.6
0.4	51.7	57.2	57.0	70.5	50.0	74.2
0.5	51.6	56.7	55.1	71.1	52.1	73.9
0.6	57.0	56.3	58.9	69.8	52.6	72.7
0.7	54.7	52.2	56.9	71.1	52.7	72.5
0.8	50.4	53.1	55.4	69.3	52.5	73.8
0.9	53.5	54.2	55.7	72.1	53.0	75.3
1.0	53.4	53.9	56.4	70.7	52.5	74.0

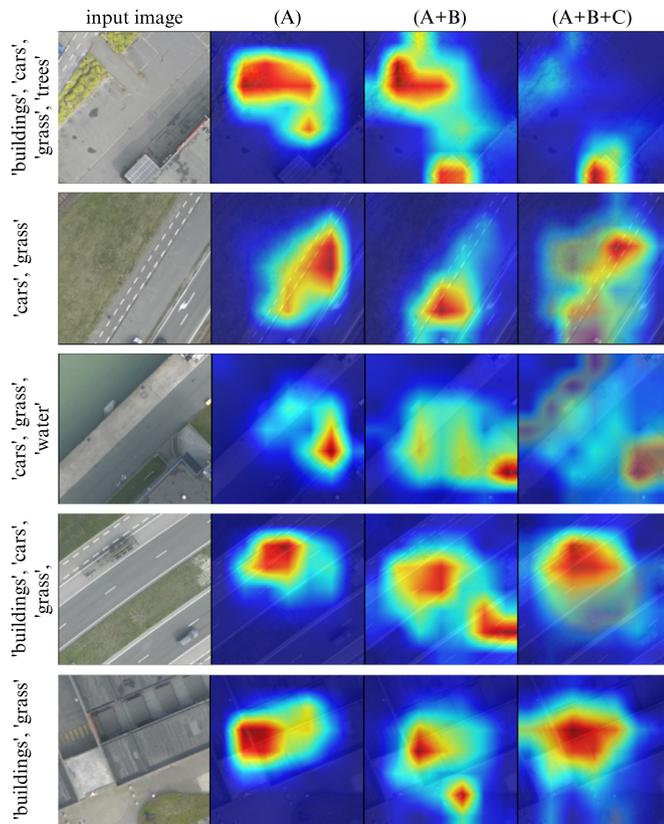


Fig. 8. Failure cases: Grad-CAM visualization for AID → DFC with ground truth multi-labels showcasing performance degradation with adaptive graphs.

multiple domains. For that purpose, an adversarial domain adaptation strategy is employed. The results performed for both single and cross-domain support the effectiveness of the proposed method in terms of model performance and size as compared to recent state-of-the-art methods. The current work is restricted to scenarios where the source and target data have shared object categories. Nevertheless, in future works, we intend to investigate a more challenging problem, namely, open set domain adaptation where only a few categories of labels are shared between source and target data.

CRedit authorship contribution statement

Inder Pal Singh: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Enjie Ghorbel:** Writing – review & editing, Supervision, Project administration, Investigation, Formal analysis. **Oyebade Oyedotun:** Conceptualization. **Djamila Aouada:** Writing – review & editing, Validation, Supervision, Resources, Project administration, Funding acquisition, Formal analysis.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Inder Pal Singh reports financial support, administrative support, article publishing charges, equipment, drugs, or supplies, statistical analysis, travel, and writing assistance were provided by University of Luxembourg. Djamila Aouada reports financial support was provided by National Research Fund (Luxembourg). One of the co-authors, Oyebade Oyedotun, is currently working with Spire Global, Luxembourg.

Data availability

Data will be made available on request.

Acknowledgments

This research was partially funded by the Luxembourg National Research Fund (FNR), grant references BRIDGES2020/IS/14755859/MEET-A/Aouada. For the purpose of open access, and in fulfillment of the obligations arising from the grant agreement, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

References

- Cai, Y., Ge, L., Liu, J., Cai, J., Cham, T.-J., Yuan, J., Thalmann, N.M., 2019. Exploiting spatial-temporal relationships for 3d pose estimation via graph convolutional networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2272–2281.
- Chaudhuri, B., Demir, B., Chaudhuri, S., Bruzzone, L., 2018. Multilabel remote sensing image retrieval using a semisupervised graph-theoretic method. *IEEE Trans. Geosci. Remote Sens.* 56 (2), 1144.
- Chen, L., Chen, H., Wei, Z., Jin, X., Tan, X., Jin, Y., Chen, E., 2022. Reusing the task-specific classifier as a discriminator: Discriminator-free adversarial domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7181–7190.
- Chen, T., Lin, L., Hui, X., Chen, R., Wu, H., 2020. Knowledge-guided multi-label few-shot learning for general image recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Chen, Z.M., Wei, X.S., Wang, P., Guo, Y., 2019b. Multi-label image recognition with graph convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5177–5186.
- Chen, T., Xu, M., Hui, X., Wu, H., Lin, L., 2019a. Learning semantic-specific graph representation for multi-label image recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 522–531.
- Cheng, X., Lin, H., Wu, X., Shen, D., Yang, F., Liu, H., Shi, N., 2022. Mltr: Multi-label classification with transformer. In: 2022 IEEE International Conference on Multimedia and Expo. ICME, IEEE, pp. 1–6.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. IEEE Conference on Computer Vision and Pattern Recognition. A Large-Scale Hierarchical Image Database, *Imagenet*, pp. 248–255.
- Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A., 2010. The pascal visual object classes (voc) challenge. *Int. J. Comput. Vis.* 88 (2), 303–338.
- Ganin, Y., Lempitsky, V., 2015. June. Unsupervised domain adaptation by back-propagation. In: International Conference on Machine Learning. PMLR, pp. 1180–1189.

- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V., 2016. Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* 17 (59), 1–35.
- Gao, B.-B., Zhou, H.-Y., 2021. Learning to discover multi-class attentional regions for multi-label image recognition. *IEEE Trans. Image Process.* 30, 5920–5932.
- Ge, W., Yang, S., Yu, Y., 2018. Multi-evidence filtering and fusion for multi-label classification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. Object Detection and Semantic Segmentation Based on Weakly Supervised Learning*, pp. 1277–1286.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Hua, Y., Mou, L., Zhu, X.X., 2019. Recurrently exploring class-wise attention in a hybrid convolutional and bidirectional LSTM network for multi-label aerial image classification. *ISPRS J. Photogramm. Remote Sens.* 149.
- Hua, Y., Mou, L., Zhu, X.X., 2020. Relation network for multilabel aerial image classification. *IEEE Trans. Geosci. Remote Sens.* 58 (7), 4558–4572.
- Inoue, N., Furuta, R., Yamasaki, T., Aizawa, K., 2018. Cross-domain weakly-supervised object detection through progressive domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5001–5009.
- Jin, W., Derr, T., Wang, Y., Ma, Y., Liu, Z., Tang, J., 2021. Node similarity preserving graph convolutional networks. In: *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. pp. 148–156.
- Kipf, T.N., Welling, M., 2017. Semi-Supervised Classification with Graph Convolutional Networks. *ICLR*.
- Krizhevsky, A., Sutskever, I., Hinton, G., 2012. Imagenet classification with deep convolutional neural networks. In: *Proc. Neural Inf. Process. Syst. Vol. 1106*.
- Lanchantin, J., Wang, T., Ordonez, V., Qi, Y., 2021. General multi-label image classification with transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16478–16488.
- Li, Y., Huang, C., Loy, C.C., Tang, X., 2016. Human attribute recognition by deep hierarchical contexts. In: *European Conference on Computer Vision*. Cham, pp. 684–700.
- Li, G., Ji, Z., Chang, Y., Li, S., Qu, X., Cao, D., 2021. ML-ANet: A transfer learning approach using adaptation network for multi-label image classification in autonomous driving. *Chin. J. Mech. Eng.* 34 (1), 1–11.
- Li, Q., Peng, X., Qiao, Y., Peng, Q., 2019. Learning category correlations for multi-label image recognition with graph networks. *arXiv preprint arXiv:1909.13005*.
- Li, M., Zhai, Y.M., Luo, Y.W., Ge, P.F., Ren, C.X., 2020. Enhanced transport distance for unsupervised domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13936–13944.
- Lin, D., Lin, J., Zhao, L., Wang, Z.J., Chen, Z., 2021. Multilabel aerial image classification with unsupervised domain adaptation. *IEEE Trans. Geosci. Remote Sens.*
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., . . . Zitnick, C.L., 2014. Microsoft coco: common objects in context. In: *European Conference on Computer Vision*. Cham, pp. 740–755.
- Liu, R., Huang, J., Li, T.H., Li, G., 2022. Causality compensated attention for contextual biased visual recognition. In: *The Eleventh International Conference on Learning Representations*.
- Long, M., Cao, Z., Wang, J., Jordan, M.I., 2018. Conditional adversarial domain adaptation. *Adv. Neural Inf. Process. Syst.* 31.
- Long, M., Zhu, H., Wang, J., Jordan, M.I., 2017. Deep transfer learning with joint adaptation networks. In: *International Conference on Machine Learning*. PMLR, pp. 2208–2217.
- Mirza, M., Osindero, S., 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Papadopoulos, K., Ghorbel, E., Aouada, D., Ottersten, B., 2020a. 25th International Conference on Pattern Recognition. *ICPR*, pp. 452–458, Vertex feature encoding and hierarchical temporal modeling in a spatio-temporal graph convolutional network for action recognition.
- Papadopoulos, K., Ghorbel, E., Oyedotun, O., Aouada, D., Ottersten, B., 2020b. Deepvi: A novel framework for learning deep view-invariant human action representations using a single rgb camera. In: *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition. FG 2020, IEEE*, pp. 138–145.
- Pennington, J., Socher, R., Manning, C.D., 2014. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing. *EMNLP, Global vectors for word representation*, *Glove*, pp. 1532–1543.
- Pham, D.D., Koesnadi, S.M., Dovletov, G., Pauli, J., 2021. In: *Adversarial, U. (Ed.), IEEE 18th International Symposium on Biomedical Imaging. ISBI, Domain Adaptation for Multi-Label Classification of Chest X-Ray*, pp. 1236–1240.
- Qu, X., Che, H., Huang, J., Xu, L., Zheng, X., 2023. Multi-layered semantic representation network for multi-label image classification. *Int. J. Mach. Learn. Cybern.* 1–9.
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al., 2021. Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning*. PMLR, pp. 8748–8763.
- Razavian, S., Azizpour, A., Sullivan, H., 2014. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, an Astounding Baseline for Recognition, and Carlsson, S., CNN features off-the-shelf, pp. 806–813, J.
- Ridnik, T., Ben-Baruch, E., Zamir, N., Noy, A., Friedman, I., Protter, M., Zelnik-Manor, L., 2021a. Asymmetric loss for multi-label classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 82–91.
- Ridnik, T., Lawen, H., Noy, A., Baruch, B., Sharir, E., 2021b. In: *trsnets: high performance gpu-dedicated architecture*. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, and Friedman, I. pp. 1400–1409, G.
- Ridnik, T., Sharir, G., Ben-Cohen, A., Ben-Baruch, E., Noy, A., 2023. MI-decoder: Scalable and versatile classification head. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 32–41.
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 618–626.
- Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y., 2014. Overfeat: Integrated Recognition, Localization and Detection Using Convolutional Networks. *ICLR*.
- Shao, J., Kang, K., Loy, C., Wang, C., 2015. In: *Deeply, X. (Ed.), Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Learned Attributes for Crowded Scene Understanding*. pp. 4657–4666.
- Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: *Bengio, Y., LeCun, Y. (Eds.), 3rd International Conference on Learning Representations. ICLR 2015*.
- Singh, I.P., Ghorbel, E., Kacem, A., Rathinam, A., Aouada, D., 2024. Discriminator-free unsupervised domain adaptation for multi-label image classification. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 3936–3945.
- Singh, I.P., Ghorbel, E., Oyedotun, O., Aouada, D., 2022a. Multi label image classification using adaptive graph convolutional networks (ML-AGCN). In: *IEEE International Conference on Image Processing*.
- Singh, I., Meiri, N., Nygyen, V., Ghorbel, D., 2023. Multi-type deepfake detection. In: *MMSp*.
- Singh, I.P., Oyedotun, O., Ghorbel, E., Aouada, D., 2022b. IML-GCN: Improved multi-label graph convolutional network for efficient yet precise image classification. In: *AAAI-22 Workshop Program-Deep Learning on Graphs: Methods and Applications*.
- Sun, X., Hu, P., Saenko, K., 2022b. Dualcoop: Fast adaptation to multi-label recognition with limited annotations. *Adv. Neural Inf. Process. Syst.* 35, 30569–30582.
- Sun, D., Ma, L., Ding, Z., Luo, B., 2022a. An attention-driven multi-label image classification with semantic embedding and graph convolutional networks. *Cogn. Comput.* 1–12.
- Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y., 2017. Graph attention networks. *Statistics* 1050.
- Wang, Y., He, D., Li, F., Long, X., Zhou, Z., Ma, J., Wen, S., 2020a. Proceedings of the AAAI Conference on Artificial Intelligence, Multi-Label Classification with Label Graph Superimposing. *Vol. 34*, pp. 12265–12272, No. 07.
- Wang, Y., Xie, Y., Fan, L., Hu, G., 2022a. STMG: Swin transformer for multi-label image recognition with graph convolution network. *Neural Comput. Appl.* 34 (12), 10051–10063.
- Wang, Y., Xie, Y., Liu, Y., Zhou, K., Li, X., 2020b. Fast graph convolution network based multi-label image recognition via cross-modal fusion. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. pp. 1575–1584.
- Wang, Y., Xie, Y., Zeng, J., Wang, H., Fan, L., Song, Y., 2022b. Cross-modal fusion for multi-label image classification with attention mechanism. *Comput. Electr. Eng.* 101, 108002.
- Wang, J., Yang, Y., Mao, J., Huang, Z., Huang, C., Xu, W., 2016. Cnn-rnn: A unified framework for multi-label image classification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2285–2294.
- Wei, Y., Xia, W., Lin, M., Huang, J., Ni, B., Dong, J., Zhao, Y., Yan, S., 2015. HCP: A flexible CNN framework for multi-label image classification. *IEEE Trans. Pattern Anal. Mach. Intell.* 38 (9), 1901–1907.
- Xia, G.S., Hu, J., Hu, F., Shi, B., Bai, X., Zhong, Y., Zhang, L., Lu, X., 2017. AID: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Trans. Geosci. Remote Sens.* 55 (7), 3965–3981.
- Yang, Y., Newsam, S., 2010. Bag-of-visual-words and spatial extensions for land-use classification. In: *International Conference on Advances in Geographic Information Systems. SIGSPATIAL*.
- Zhang, Y., Liu, T., Long, M., Jordan, M., 2019. Bridging theory and algorithm for domain adaptation. In: *International Conference on Machine Learning*. PMLR, pp. 7404–7413.
- Zhu, F., Li, H., Ouyang, W., Yu, N., Wang, X., 2017. Learning spatial regularization with image-level supervisions for multi-label image classification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5513–5522.