











GKNet: Group K-Nearest Neighbor based Graph Convolutional Network for Multi-Label Image Recognition

Ruijie Yao^{1,4} , Sheng Jin^{2,4} , Lumin Xu³ , Wang Zeng⁴ , Wentao Liu⁴ ,
Chen Qian^{1,4}  , Ping Luo^{2,5} , and Ji Wu¹  


¹ Tsinghua University ² The University of Hong Kong ³ The Chinese University of Hong Kong ⁴ SenseTime Research and Tetras.AI ⁵ Shanghai AI Laboratory
{yrj21@mails, qianc18@mails, wuji_ee@mail}.tsinghua.edu.cn

Abstract. Multi-Label Image Recognition (MLIR) is a challenging task that aims to predict multiple object labels in a single image while modeling the complex relationships between labels and image regions. Although convolutional neural networks and vision transformers have succeeded in processing images as regular grids of pixels or patches, these representations are sub-optimal for capturing irregular and discontinuous regions of interest. In this work, we present the first fully graph convolutional model, Group K-nearest neighbor based Graph convolutional Network (GKNet), which models the connections between semantic label embeddings and image patches in a flexible and unified graph structure. To address the scale variance of different objects and to capture information from multiple perspectives, we propose the Group KGCN module for dynamic graph construction and message passing. Our experiments demonstrate that GKNet achieves state-of-the-art performance with significantly lower computational costs on the challenging multi-label datasets, *i.e.* MS-COCO and VOC2007 datasets. Codes are available at <https://github.com/jin-s13/GKNet>.

1 Introduction

Multi-label image recognition (MLIR) (also referred to as multi-label classification) is a fundamental task in computer vision, which aims to predict a set of labels of a single image. Compared with single-label image recognition, MLIR is more challenging due to its combinatorial nature. It has received increasing attention because of its broad real-world applications, such as human attribute recognition [19] and scene understanding [26].

Modeling the relationships between each target label and the corresponding image regions is critical yet challenging for MLIR. The regions that relate to a certain label may be complicated and even discontinuous. For example, to recognize the presence of dogs in an image, multiple regions need to follow

 : Corresponding authors.




Fig. 1: Illustration of feature extraction in CNN, vision transformer, and graph convolutional network (GCN). (a) CNN excels at processing continuous regions but struggles with irregular regions of interest. (b) Vision transformers handle complex regions of interest but introduce redundant interference from the background. (c) GCN constructs connections between the destination node and multiple objects of interest distributed in different spatial locations.

with interest, since there may exist multiple dogs. Convolutional neural networks (CNN) [12, 14, 34] regard the image as grids of pixels and apply sliding convolutional kernels according to spatial positions. As shown in Fig. 1 (a), CNN process continuous regions well but struggles to fit irregular regions of interest. In contrast, vision transformers [6] treat the image as a sequence of patches and extract the visual features from the patches via attention. As shown in Fig. 1 (b), vision transformers can handle complicated regions of interest by capturing the whole image features. However, interactions with the background increase costs and affect performance. For small-size target objects, most patches belong to the background thus the sum of their attention scores cannot be ignored. Graph methods like Vision GNN [11] views visual patches as nodes, forming a flexible approach to extract features by representing images as graphs. As shown in Fig. 1 (c), the graph representation constructs the connections according to the semantic meanings and focuses on the regions of interest distributed in different spatial locations.

In this work, we propose the first fully graph convolutional network (GCN) for the task of MLIR, termed **Group K**-nearest neighbor based **Graph** convolutional **Network** (GKGNet). GKGNet regards both the image patches and target labels as graph nodes and processes them in a unified graph structure. GKGNet constructs two distinct graphs: a cross-level graph that models the label-object relationship between target labels and image patches, and a patch-level graph that processes and updates the image features between patches. Thus, GKGNet enables the adaptive integration of features from patches of interest, even in the presence of irregular and discontinuous regions, leading to the effective updating of the unified graph representations of visual features and label embeddings.

Graph construction is the key to the success of GCN. The K -nearest neighbor (KNN) graph, where there is an edge between each node and its K most similar nodes, is one of the most widely used graph structures [11]. However, this approach is sub-optimal for the following reasons. First, the number of neighbors K controls the size of the area where region features will be extracted and

aggregated. Large K will lead to feature over-smoothing and involve interference of the invalid background, while small K will affect feature extraction and message passing. A fixed K cannot adaptively handle objects with different scales. Second, this method finds the neighbors by measuring the distance between the destination node and the candidate source nodes. It is difficult to define a unique measurement of distance to represent the rich dimensions of “high-level” labels.

Therefore, we propose the Group KNN-based GCN (Group KGCN) module, which splits node features into multiple groups and constructs connections among each group. By designing each group to potentially connect to different source nodes, the Group KNN approach allows a destination node to interact with a dynamic number of source nodes and handle objects with different scales. Also, this dynamic connectivity enables the model to capture information from multiple perspectives, which is crucial for MLIR. For instance, recognizing the class “zebra” requires consideration of both shape and appearance features. Furthermore, MLIR involves capturing multi-label correlations, where the context surrounding the target objects is also important for accurate recognition [2, 4, 29–32]. Therefore, by allowing for flexible and varied connectivity patterns, the Group KNN approach is able to model these complex relationships between image regions and labels.

Experiments on two well-known benchmarks, *i.e.* MS-COCO and VOC2007, verify the effectiveness of the proposed method. Our method achieves new state-of-the-art performance with significantly lower computational costs.

This paper’s contribution can be summarized as follows:

- We introduce GKGNet, the first fully graph convolutional model for multi-label image recognition, which builds unified graph representations for visual features and label embeddings, explicitly modeling the flexible relationship between labels and regions of interest.
- To handle scale variance and capture diverse perspectives, we propose the Group KGCN module for dynamic graph construction and message passing.
- Our proposed model is validated on popular benchmark datasets, including MS-COCO and VOC2007. It consistently outperforms previous state-of-the-art approaches with lower computational complexity.

2 Related Work

2.1 Multi-Label Image Recognition

Multi-label image recognition (MLIR) (also referred to as multi-label classification) is an extension of image classification, which aims to predict multiple class labels for an image. Several approaches [12, 14, 34] design powerful Convolutional Neural Network (CNN) architectures for the task of MLIR. CNN extracts the visual features of the input image and predicts multiples labels according to the same features. Instead of explicitly capturing the regions of interest, these methods treat each label equally and simply consider MLIR as multiple binary




Fig. 2: Overview of GKGNet. GKGNet splits the input image into a set of patch nodes, and regards the learnable label embeddings as label nodes. Four-stage network is applied to process the patch nodes and label nodes in the unified graph structure. The number of patch nodes is reduced after each stage to extract multi-scale visual features. At each stage, the patch nodes are first updated via Patch-Level Group KGCN modules, and then Cross-Level Group KGCN modules updates the label nodes by building the connections between target labels and image regions of interest. The output patch nodes and label nodes of the last stage are combined for multi-label prediction.

classification tasks. Recently, some works [6, 17, 22] explore to use transformers for the task of MLIR to capture semantic relationships among the labels. C-Tran [17] uses a transformer encoder to process image features and label embeddings simultaneously, while Q2L [22] employs a transformer decoder to explicitly exchange information between target labels and image features. These methods predict labels using visual features of the whole image, leading to redundant background interference and performance hindrance. Moreover, such designs demand high computational costs, particularly as image resolution increases.

There are several works [1, 4, 5, 29, 30, 32] applying graph convolutional network (GCN) for MLIR. They use CNN for image representation learning and only use GCN to model the relationship between multiple labels. However, these methods do not explore the important connections between the global classification label and spatial visual regions. Moreover, the image representations learned by CNN are not well-aligned with the semantic embeddings of labels processed by GCN, which hinders message passing. To deal with the above problems, we instead propose a fully graph convolutional network, termed GKGNet, to capture the relationship between target labels and the spatial areas in a unified graph representation. GKGNet adaptly searches the regions of interest for each target label in the flexible graph structure. It is able to fit in irregular areas and capture complex correlations without introducing any computational cost.

2.2 Graph Convolutional Network

Graph is a versatile and flexible data structure and it can process any kind of data that can be converted into a set of nodes and edges. They can represent both non-Euclidean data, such as social networks, and Euclidean data, such as images [33]. Graph convolutional networks (GCN) [13] are introduced to learn

graph representations using convolutional operations on the graph’s Laplacian matrix based on spectral graph theory. GCN is highly effective for message passing and correlation modeling in computer vision tasks, such as multi-label classification [4], relationship proposal generation [28], scene graph generation [35], and human pose estimation [15]. Recently, a GCN-based backbone network called ViG [11] has been proposed to directly convert images into graph structures to learn visual representations. This approach divides an image into patches and considers them as nodes in a graph. Edges are constructed by identifying K-nearest neighbors (KNN) for each node, and effective message passing captures the spatial relationships between visual patches. However, the KNN-based approach limits the destination node to aggregate information from a fixed number of source nodes, which may not be appropriate for tasks involving target labels with regions of different sizes and shapes. To address this limitation, we propose the Group KGCN module, which can adapt to different numbers of neighbors based on practical requirements.

3 Method

3.1 Overview

We propose a fully graph convolutional network (GCN), termed Group K-nearest neighbor based Graph convolutional Network (GKNet), for the task of multi-label image recognition. The overview of GKNet is illustrated in Fig. 2. Inspired by ViG [11], the input image is divided into N image patches, and each patch is transformed into a visual feature vector with dimension C by a fully-connected layer. These feature vectors are represented as patch nodes in our graph representations. The learnable label embeddings are viewed as label nodes, which have the same feature dimension C as the patch nodes. The patch nodes and label nodes are processed by four hierarchical stages, and the number of patch nodes is reduced after each stage to extract multi-scale visual features. Each stage consists of two kinds of Group KGCN modules, *i.e.* Patch-Level Group KGCN module and Cross-Level Group KGCN module. Patch-Level Group KGCN module serves as image feature extractor to capture the spatial relationship between visual features, and Cross-Level Group KGCN module establishes the cross-correlation between target labels and the related regions of interest. In order to construct flexible connections between nodes, Group KGCN module allows each node to correlate with different number of patch nodes according to practical needs. Finally, the classifier is applied to the output patch nodes and label nodes, and both of them are taken into consideration for class possibility prediction.

3.2 Group KGCN Module

Our GKNet paradigm entails the representation of both patch and label features in a graph form, wherein edges are established among neighboring nodes to facilitate message passing. A conventional technique for graph construction

involves identifying the K-nearest neighbors (KNN) of the destination node, as depicted in Fig. 3 (a), and subsequently updating the destination node’s features based on those of its neighbors, *i.e.* source nodes. However, such a methodology restricts itself to a fixed number of source nodes and may not be optimal for the task of MLIR, because the target label may be associated with regions of varying sizes and different number of source nodes should be connected. To address this issue, we propose Group K-nearest neighbors (Group KNN), which provides greater flexibility in constructing the graph structure, for Group KNN-based GCN (Group KGCN) module, enabling effective message passing.

Group KNN. Group KNN based graph construction divides both source nodes $\mathbb{S} = \{S_1, S_2, \dots, S_{N_S}\} \in \mathbb{R}^{N_S \times C}$ and destination nodes $\mathbb{D} = \{D_1, D_2, \dots, D_{N_D}\} \in \mathbb{R}^{N_D \times C}$ into G groups of sub source nodes and sub destination nodes by splitting feature dimensions, where C denotes the dimension of the destination node and source node, and N_S and N_D are numbers of source nodes and destination nodes respectively. Each sub destination node selects to connect with sub source nodes within the same group. For group g , the i -th sub destination node $D_{ig} \in \mathbb{R}^{\frac{C}{G}}$ ($i \in [1, N_D], g \in [1, G]$) searches for the K-nearest neighbors among the corresponding sub source nodes $\mathbb{S}_g = \{S_{1g}, S_{2g}, \dots, S_{N_Sg}\} \in \mathbb{R}^{N_S \times \frac{C}{G}}$. Cosine similarity is used for neighborhood search.

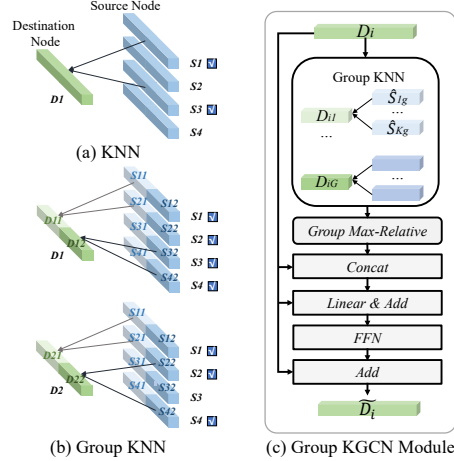


Fig. 3: Illustration of Group KGCN. (a) Traditional KNN based graph construction (K=2). (b) Group KNN based graph construction (G=2, K=2). The blue check marks indicate the source nodes are selected. (c) Structure of Group KGCN module.

The utilization of distinct groups allows the destination node to establish connections with a variable number of source nodes, ranging from K to $K \times G$. As shown in Figure 3 (b), the sub destination node D_{11} links to the sub source

nodes S_{11} and S_{21} , while the sub destination node D_{12} connects with the sub source nodes S_{31} and S_{41} . This enables the destination node D_1 to interact with four source nodes: S_1 , S_2 , S_3 , and S_4 . This scenario typically occurs when the destination node concerns a broader area of interest, such as a large target object. Conversely, when the neighbors of distinct groups overlap, the number of selected source nodes decreases. For instance, both sub destination nodes of D_2 connect with the source nodes S_2 , resulting in an interaction with only three source nodes: S_1 , S_2 , and S_4 . This scenario typically occurs when the destination node involves only small regions of interest, such as a small target object, thereby helping to circumvent irrelevant information during message passing.

Module structure. As shown in Fig. 3 (c), the Group KGCN module conducts message passing between the destination node and source nodes after the graph is constructed via Group KNN. The sub destination node D_{ig} is updated according to its K-nearest neighbor source nodes $\{\hat{S}_{1g}, \hat{S}_{2g}, \dots, \hat{S}_{Kg}\} \subseteq \mathbb{S}_g$ via group max-relative graph convolution. \hat{S} represents the selected nearest neighbors.

Group max-relative graph convolution is an extension of [18]. For each sub destination node, it aggregates features of its group-wise neighbors by maximizing every dimension of the relative feature $D_{ig} - \hat{S}_{kg}$ as follows:

$$D'_{ig} = \max(\{D_{ig} - \hat{S}_{kg} | k \in [1, K]\}). \quad (1)$$

Then all the updated sub destination nodes are processed to output the renewed destination node \tilde{D}_i as follows:

$$\tilde{D}_i = D_i + FFN(D_i + \text{Linear}(\text{Concat}(D_i, \{D'_{ig}\}))), \quad (2)$$

where *Concat* indicates the concatenation operation, and *Linear* is a fully connected layer. *FFN* denotes the feed-forward network. The original destination node D_i and its G sub-destination nodes processed by group max-relative graph convolution are concatenated together for destination node feature update. The residual structure is used to relieve the over-smoothing problem.

Computation cost analysis. Our Group KGCN module improves model accuracy without adding computational costs. Traditional KNN-based GCNs have a distance calculation cost of $O(N_S \times N_D \times C)$ for N_S source nodes, N_D destination nodes, and dimension C . The Group KGCN module has a cost of $O(G \times N_S \times N_D \times \frac{C}{G}) = O(N_S \times N_D \times C)$, which is the same.

3.3 Patch-Level and Cross-Level Group KGCN

Each GKGNet stage consists of two types of Group KGCN modules. Patch-Level Group KGCN module updates visual features within patches and Cross-Level Group KGCN module builds connections between labels and patches.

We apply Patch-Level Group KGCN modules, where both destination and source nodes are patch nodes, to extract image features. This graph-based approach allows modeling semantic relationships among regions. Patch-Level Group KGCN modules capture the parts of an object or multiple objects of the same category distributed in different locations for more robust visual representations.

Besides the relationship among patch-level visual features, we take into consideration the cross-level label-patch correlation. We construct Cross-Level Group KGCN modules that enable message passing from patch nodes (source nodes) to label nodes (destination nodes), to learn label embeddings from the region features. The graph representation provides a flexible way to extract global context information, which are important for the task of multi-label image recognition (MLIR) where the target objects may be irregular and discontinuous. Group KNN based graph construction enables each label to capture objects of different scales and multi-label correlations. By constructing the unified graph representation for labels and patches, the cross-level graph effectively benefit label recognition from the regions of interest in the same representation space.

3.4 Classifier and Loss Function

The classifier utilizes both patch nodes and label nodes outputted by the last stage for MLIR. The final scores are:

$$Y = \text{Sigmoid}(Y_{x_p} + Y_{x_l}), \quad (3)$$

where $Y_{x_p} \in \mathbb{R}^L$ and $Y_{x_l} \in \mathbb{R}^L$ are the logits predicted by patch nodes and label nodes using the fully connected layers respectively. L is the number of target categories. Label smooth loss \mathcal{L}_{smooth} [27] and asymmetric loss \mathcal{L}_{asy} [25] are applied to supervise the training process. The total training loss is $\mathcal{L} = \mathcal{L}_{smooth} + \mathcal{L}_{asy}$.

4 Experiments

We evaluate our model on several benchmark datasets, including MS-COCO [20] and Pascal VOC [9]. We follow common practice [17, 29, 30] to report the average of Overall Recall (OR), Overall Precision (OP), Overall F1-score (OF1), per-Class Recall (CR), per-Class Precision (CP), per-Class F1-score (CF1) and the mean Average Precision (mAP) as the evaluation metrics. OF1, CF1 and mAP are emphasized as the most important metrics as OR, OP, CR and CP are easily affected by the classification threshold. We set the threshold as 0.5 for recall, precision, and F1 score in our experiments.

4.1 Implementation Details

Numbers of Patch-Level Group KGCN modules are set as [2, 2, 6, 2], and numbers of Cross-Level Group KGCN modules are [1, 1, 1, 1] for four hierarchical stages

Table 1: Comparisons with state-of-the-art methods on MS-COCO. All the methods adopt models pre-trained on ImageNet-1K dataset. [†] means using model EMA. We report multiple evaluation metrics (higher is better), among which mAP, CF1, and OF1 are the primary ones. GKGNet significantly outperforms the existing approaches in terms of both accuracy and efficiency.

	Resolution	Param(M)	FLOPs(G)	All							Top3	
				mAP	CP	CR	CF1	OP	OR	OF1	CF1	OF1
ResNet-101 [12]	224 × 224	44.5	7.8	78.3	80.2	66.7	72.8	83.9	70.8	76.8	69.7	73.6
SRN [34]	224 × 224	76.8	9.0	77.1	81.6	65.4	71.2	82.7	69.9	75.8	67.4	72.9
IDA-SwinS(H) [21]	224 × 224	48.9	8.7	80.6	64.5	81.1	71.8	65.8	83.9	73.8	71.4	74.5
MLtr [6]	224 × 224	33.0	-	81.9	80.7	71.5	75.2	81.4	76.3	78.1	-	-
GKGNet (Ours)	224 × 224	33.3	5.2	82.0	81.7	73.1	77.1	83.1	76.3	79.6	73.7	76.1
CADM [3]	448 × 448	-	-	82.3	82.5	72.2	77.0	84.0	75.6	79.6	73.5	76.0
ML-GCN [4]	448 × 448	44.9	31.5	83.0	85.1	72.0	78.0	85.8	75.4	80.3	74.6	76.7
KSSNet [29]	448 × 448	173.8	-	83.7	84.6	73.2	77.2	87.8	76.2	81.5	-	-
MS-CMA [31]	448 × 448	-	-	83.8	82.9	74.4	78.4	84.4	77.9	81.0	74.9	77.1
MCAR [10]	448 × 448	-	-	83.8	85.0	72.1	78.0	88.0	73.9	80.3	75.1	76.7
TDRG [32]	448 × 448	68.3	42.2	84.6	86.0	73.1	79.0	86.6	76.4	81.2	75.0	77.2
Q2L-R101 [22] [†]	448 × 448	193.6	51.4	84.9	84.8	74.5	79.3	86.6	76.9	81.5	73.3	75.4
IDA-SwinS(H) [21]	448 × 448	48.9	35.0	85.5	68.9	85.6	76.3	69.8	87.5	77.7	75.2	77.4
GKGNet (Ours)	448 × 448	34.0	21.9	86.7	86.4	77.1	81.5	87.3	79.7	83.3	77.0	78.8
SSGRL [2]	576 × 576	92.3	68.5	83.8	91.9	62.5	72.7	93.8	64.1	76.2	76.8	79.7
KGGR [1]	576 × 576	-	-	84.3	85.6	72.7	78.6	87.1	75.6	80.9	75.0	77.0
MCAR [10]	576 × 576	-	-	84.5	84.3	73.9	78.7	86.9	76.1	81.1	75.3	77.0
ADD-GCN [30]	576 × 576	48.2	52.7	85.2	84.7	75.9	80.1	84.9	79.4	82.0	75.8	77.9
C-Tran [17]	576 × 576	120.4	84.2	85.1	86.3	74.3	79.9	87.7	76.5	81.7	76.0	77.6
TDRG [32]	576 × 576	68.3	69.8	86.0	87.0	74.7	80.4	87.5	77.9	82.4	76.2	78.1
IDA-R101(H) [21]	576 × 576	53.2	54.1	86.3	-	-	80.4	-	-	82.5	76.4	78.2
IDA-SwinS(H) [21]	576 × 576	48.9	64.0	86.4	69.6	86.5	77.1	69.6	88.8	78.1	75.6	77.8
Q2L-R101 [22] [†]	576 × 576	193.6	80.8	86.5	85.8	76.7	81.0	87.0	78.9	82.8	76.5	78.3
GKGNet (Ours)	576 × 576	34.7	40.1	87.7	87.0	78.5	82.5	87.6	81.0	84.2	77.6	79.3

respectively. The feature dimensions C for each stage are set as [80, 160, 400, 640]. We choose $K = 9$ for the number of neighbors and $G = 2$ for the number of groups in all the Group KGCN modules. ImageNet-1K [8] pre-trained weights of Pyramid ViG-S [11] are used for Patch-Level Group KGCN modules initialization, while Cross-Level Group KGCN modules and learnable label embeddings are randomly initialized. During training, we utilize AdamW optimizer [23] with a weight decay rate of 0.05 and momentum of 0.9 to train GKGNet for 80 epochs. The initial learning rate is 1×10^{-4} with a linear warm-up of 1×10^{-3} . We reduce the learning rate by a factor of 0.1 at the 5th and 50th epoch for Pascal VOC dataset, and at the 10th and 50th epoch for MS-COCO [20] dataset respectively. All experiments are based on MMClassification [7].

4.2 Comparisons on MS-COCO Dataset

MS-COCO [20] is one of the most widely used benchmark datasets for multi-label classification. It contains 122,218 images that are composed of 82,081 training images and 40,137 validation images. There exist 80 kinds of object labels in total, with 2.9 labels on average for each image. We report results on three commonly used input resolution settings on the dataset: 224×224 , 448×448 , and 576×576 . All methods apply the models pre-trained on ImageNet-1K [8] as the common practice. For our GKGNet, only Patch-Level Group KGCN uses pre-

trained weights, while other parts use random initialization. For fair comparisons, we do not report results that apply much larger models (*e.g.* Swin Large with 197M parameters) and extensive pre-training data (*e.g.* ImageNet21K), which involves a different experiment setting from ours.

As shown in Table 1, our method achieves the best performance on all the input resolution settings in terms of mAP, CF1, and OF1. Compared with transformer based methods (*e.g.* C-Tran [17] and Q2L [22]), the flexibility of graph structure enables the model to focus only on the regions of interest and to avoid background interference, which saves computational cost and improves the performance. For example, GKNet outperforms Q2L by 1.2% mAP with half the FLOPs under the 576×576 resolution setting. Kindly note that Q2L also utilizes Exponential Moving Average (EMA) technique that maintains moving averages of the trained model parameters to improve training robustness and model performance. ML-GCN [4], KSSNet [29], MS-CMA [31], ADD-GCN [30], SSGRL [2], and TDRG [32] are GCN based models. We find that our proposed GKNet outperforms them by a large margin, which demonstrates the superiority of the unified design of patch-level and cross-level graph-based interaction, and the proposed Group KGCN modules that construct the graph structure adaptively according to the scale of target objects. Under the 448×448 resolution setting, our model improves upon ML-GCN, KSSNet, MS-CMA, and TDRG by 3.7%, 3.0%, 2.9%, and 2.1% mAP respectively. We also compare with the model using the transformer backbone (IDA-SwinS(H) [21]). Our method shows consistent superiority over the transformer backbone at various resolutions. For example, under the 448×448 resolution setting, GKNet outperforms IDA-SwinS(H) (86.7 mAP v.s. 85.5 mAP).

Table 2: We compare GKNet with state-of-the-art methods using the same feature extractor ViG on MS-COCO (448×448 input size) [†] means using model EMA.

	All			Top3	
	mAP	CF1	OF1	CF1	OF1
ViG [11]	82.5	76.6	80.0	72.7	76.2
+ C-Tran [17]	84.5	79.2	80.8	75.6	77.0
+ Q2L [22]	85.5	77.1	78.2	74.9	76.8
+ Q2L [22] [†]	85.9	79.2	80.8	76.1	77.9
GKNet	86.7	81.5	83.3	77.0	78.8

Same backbone. To decouple the effect of image feature extractor, we apply ViG to state-of-the-art methods (C-Tran and Q2L) using official codes and recommended training settings. The result of ViG is also reported as the baseline by replacing the final linear layer as [12]. The same ImageNet-1K pre-trained weights are applied to all the models. The results of Q2L with and without the Exponential Moving Average (EMA) technique are both reported. As shown in Table 2, our proposed GKNet improves upon ViG baseline by 4.2% mAP,

Table 3: Comparisons with state-of-the-art methods on Pascal VOC2007 dataset. We report the average precision in each category, and the mean average precision (mAP) of all the categories. All the models are pre-trained on the MS-COCO (576×576 input size). Our proposed GKGNet outperforms the previous state-of-the-arts.

	Aero	Bike	Bird	Boat	Bottle	Bus	Car	Cat	Chair	Cow	Table	Dog	Horse	Mbike	Person	Plant	Sheep	Sofa	Train	TV	mAP
SSGRL	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	97.0	99.0	86.9	98.1	85.8	99.0	93.7	95.0
ASL	99.9	98.4	98.9	98.7	86.8	98.2	98.7	98.5	83.1	98.3	89.5	98.8	99.2	98.6	99.3	89.5	99.4	86.8	99.6	95.2	95.8
ADD-GCN	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
Q2L	99.9	98.9	99.0	98.4	87.7	98.6	98.8	99.1	84.5	98.3	89.2	99.2	99.2	99.3	90.2	98.8	88.3	99.5	95.5	96.1	
GKGNet (Ours)	99.9	99.4	99.2	99.4	87.0	98.2	99.1	99.6	88.4	99.5	92.6	99.5	99.5	98.7	99.5	89.5	99.3	90.4	99.7	97.2	96.8

Table 4: Effect of model components in GKGNet. The experiments are conducted on MS-COCO (448×448 input size). P, C, and G represent Patch-Level Graph, Cross-Level Graph, and Group KNN, respectively.

P	C	G	mAP	CF1	OF1
			79.9	74.6	78.7
✓			82.5	76.6	80.0
✓	✓		85.5	80.4	82.6
✓	✓	✓	86.7	81.5	83.3

Table 5: Effect of Group KNN on general classification. Top-1 accuracy of the original Pyramid ViG-Tiny and the one enhanced with our Group KNN are reported on general classification datasets (448×448 input size).

	ImageNet-1K	CIFAR-10
Pyramid ViG-Ti	78.2	94.6
+ Group KNN	79.3	94.9
	CIFAR-100	Flowers
Pyramid ViG-Ti	74.4	83.6
+ Group KNN	76.5	87.2

and outperforms C-Tran and Q2L by a clear margin, which demonstrates the effectiveness of our specific design of GKGNet for the task of MLIR.

4.3 Comparisons on Pascal VOC 2007 Dataset

Pascal VOC 2007 [9] is a commonly used benchmark dataset for MLIR, which contains 20 label categories. It has 9,963 images in total, and is divided into a train-val dataset (5,011 images) and a test dataset (4,952 images). The model is trained on the train-val dataset and evaluated on the test dataset as the common settings [2, 22, 25, 30]. For fairness, we follow the previous works [2, 30] to pre-train the model on the MS-COCO and report results at 576×576 resolution.

As shown in Table 3, our approach achieves 0.7% mAP improvement (96.8% vs 96.1% mAP) upon the previous state-of-the-art method Q2L. GKGNet consistently outperforms the previous methods on 14 of 20 categories, and achieves the competitive results on the remaining 6 categories. We also notice that the gain is more significant for these challenging categories, *e.g.* chair and table.

4.4 Ablation Study

Effect of model components. In GKGNet, the Patch-Level and Cross-Level Group KGCN modules model intricate object-object and label-object relation-

Table 6: Effect of object scales. We report mAP for varying object sizes on MS-COCO with 448×448 input size.

	Small	Medium	Large
TDRG [32]	31.3	69.4	85.1
Q2L [22]	30.7	70.2	85.6
GKGNet	35.6	73.6	86.6

Table 7: Sensitivity to random initial values. We report results on MS-COCO with 576×576 input size.

	All			Top3	
	mAP	CF1	OF1	CF1	OF1
GKGNet	87.7 ± 0.0	82.2 ± 0.1	83.8 ± 0.1	77.6 ± 0.0	79.3 ± 0.0

ships, while Group KNN adaptively constructs graphs. Table 4 validates these components individually. The first line shows the result of ResNet-50, which has comparable parameters to our model. Patch-Level graphs capture spatial relationships, achieving a 2.6% gain. The Cross-Level graphs use unified representations of patch features and label embeddings to extract specific visual features for each target label, resulting in a 3.0% gain. Group KNN based graph construction handles scale variances, enhances feature diversity, and captures label correlations, leading to a significant improvement (86.7% vs. 85.5% mAP).

Additionally, Group KNN demonstrates potential for general image recognition. We apply Group KNN to the pyramid ViG-Tiny [11], and maintain fairness with identical training settings and hyper-parameters across all datasets. Training details are available in Supplementary. As shown in Table 5, our proposed Group KNN achieves a substantial improvement in top-1 accuracy on ImageNet-1K, CIFAR-10, CIFAR-100 [16], and Flowers [24] datasets, while maintaining the same model parameters and computational complexity (10.7M parameters and 1.7B FLOPs). The significant performance gain serves as strong evidence supporting the superiority of Group KNN based graph construction.

Effect of object scales. MS-COCO divides objects into small (area $< 32^2$), medium ($32^2 < \text{area} < 96^2$) and large (area $> 96^2$). Results for various target object scales are shown in Table 6, with GKGNet achieving a notable 4.9% mAP gain for small objects compared to Q2L. Q2L faces challenges with small objects due to excessive whole-image interactions and background interference in transformers. In contrast, GKGNet selects regions of interest, saving computation and bypassing background disruptions.

Sensitivity to random initial values. We conduct experiments with three different random seeds for initialization and report their performances via mean and variance. Experiments show that the model’s performance is robust to different initial values.

Effect of G and K . For Group KGCN module, the number of groups G and the neighbor number K are important hyper-parameters. We explore the effect of G from 1 to 8 and show the results in Fig. 4 (Left). Setting $G = 2$ significantly improves the performance upon $G = 1$ (traditional KNN), which validates the importance of dynamic neighbor number choosing. Further increasing G does




Fig. 4: Effect of the number of groups G (Left) and number of neighbors K (Right).

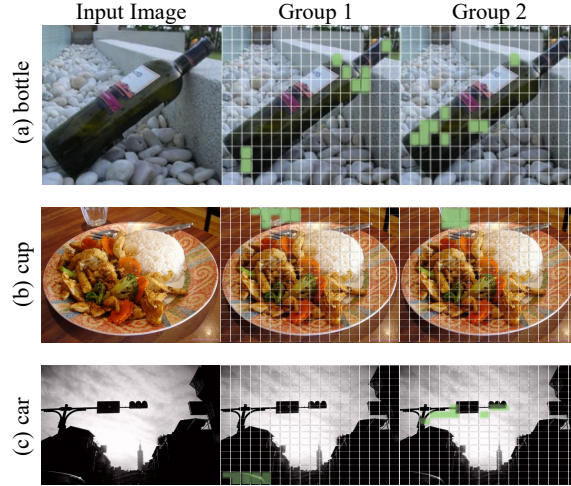


Fig. 5: Visualization of the learned connections between label node and patch nodes in the Cross-Level Group KGCN module. The colored blocks indicate that the patches are connected to the label “bottle”, “cup”, or “car”.

not bring more gains, so we choose $G = 2$ for simplicity. For neighbor number K , as shown in Fig. 4 (Right), having too few neighbors (*e.g.* $K = 3$) leads to the loss of information, while larger K performs similarly well. We set $K = 9$, aligning with ViG.

4.5 Visualization and Analysis

Cross-Level Group KGCN module. In Fig. 5, we visualize the connections of different groups at the last stage of Cross-Level Group KGCN module. Cross-Level Group KGCN module adaptively connects each label embedding to patches of different scales, shapes and locations. In Fig.5 (a), for the label “bottle”, one group focuses on the bottom and neck, and the other captures the body, sufficiently avoiding information loss. In Fig.5 (b), for the smaller “cup”, different groups select several overlapped patches, effectively capturing foreground regions while reducing background interference. In Fig. 5 (c), the Group KGCN module captures co-occurring categories through multiple perspectives. For the label “car”, the first group identifies car areas, while the second group locates the area of traffic lights, enhancing semantic relations and producing robust predictions.

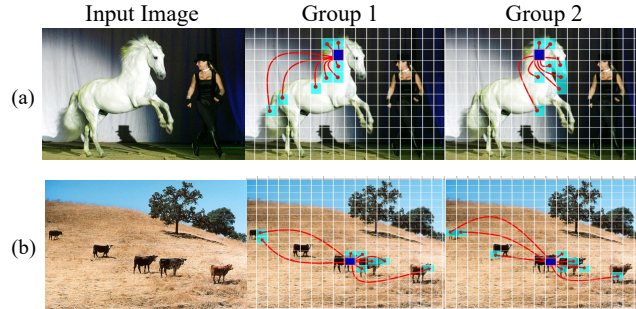


Fig. 6: Visualization of connections in the Patch-Level Group KGCN module. Deep blue represents the destination node, and baby blue patches are its selected neighbors. Red lines depict the connections between patch nodes.

Patch-Level Group KGCN module. Patch-Level Group KGCN module updates features of each patch by interacting with neighbor patches to capture semantic visual relationships flexibly. In Fig.6 (a), the destination patch (deep blue) selects two different groups of source patches (baby blue) to interact with different parts of the horse: forequarters and hindquarters. It learns message passing among semantically related patches, covering regions of interest through multiple groups. In Fig.6 (b), the destination patch selects spatially distributed foreground patches (multiple cows) as neighbors, enabling long-range correlation. The overlapped selection of two groups avoids background interference due to the small size of foreground areas.

5 Conclusion

In this paper, we propose GKGNet, a novel fully graph convolutional model for the task of MLIR. We pioneer to study the unified graph representations for both visual features and label embeddings. Group KGCN module is proposed for dynamic graph construction and message passing. It is effective in handling the scale variance of different objects, capturing information from different perspectives, and modeling the co-occurrence of different objects. Comprehensive experiments on public benchmark datasets, *i.e.* MS-COCO and VOC2007, demonstrate the effectiveness of our method. We hope the idea of unifying multi-modality features with dynamic graph representations can be broadly useful and our work can draw the community’s attention to this promising direction. In the future, we plan to extend our work to a wider range of graph-based learning problems, *e.g.* point clouds and social networks.

Limitations. Due to resource limitation, we are not able to present results of larger-scale models and extensive pre-training, *e.g.* ImageNet22K, which requires thousands of GPU hours. We’re still actively seeking collaborations and more resources to scale up both the model and data.

Acknowledgement. This paper is partially supported by the National Key R&D Program of China No.2022ZD0161000 and the General Research Fund of Hong Kong No.17200622 and 17209324.

References

1. Chen, T., Lin, L., Hui, X., Chen, R., Wu, H.: Knowledge-guided multi-label few-shot learning for general image recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* (2020)
2. Chen, T., Xu, M., Hui, X., Wu, H., Lin, L.: Learning semantic-specific graph representation for multi-label image recognition. In: *Int. Conf. Comput. Vis.* pp. 522–531 (2019)
3. Chen, Z.M., Wei, X.S., Jin, X., Guo, Y.: Multi-label image recognition with joint class-aware map disentangling and label correlation embedding. In: *Int. Conf. Multimedia and Expo.* pp. 622–627. *IEEE* (2019)
4. Chen, Z.M., Wei, X.S., Wang, P., Guo, Y.: Multi-label image recognition with graph convolutional networks. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5177–5186 (2019)
5. Chen, Z., Wei, X.S., Wang, P., Guo, Y.: Learning graph convolutional networks for multi-label recognition and applications. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
6. Cheng, X., Lin, H., Wu, X., Shen, D., Yang, F., Liu, H., Shi, N.: Mltr: Multi-label classification with transformer. In: *Int. Conf. Multimedia and Expo.* pp. 1–6. *IEEE* (2022)
7. Contributors, M.: Openmmlab’s image classification toolbox and benchmark. <https://github.com/open-mmlab/mmlclassification> (2020)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 248–255 (2009)
9. Everingham, M., Eslami, S., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. *Int. J. Comput. Vis.* **111**(1), 98–136 (2015)
10. Gao, B.B., Zhou, H.Y.: Learning to discover multi-class attentional regions for multi-label image recognition. *IEEE Trans. Image Process.* **30**, 5920–5932 (2021)
11. Han, K., Wang, Y., Guo, J., Tang, Y., Wu, E.: Vision gnn: An image is worth graph of nodes. In: *Adv. Neural Inform. Process. Syst.* (2022)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 770–778 (2016)
13. Henaff, M., Bruna, J., LeCun, Y.: Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163* (2015)
14. Jia, J., Chen, X., Huang, K.: Spatial and semantic consistency regularizations for pedestrian attribute recognition. In: *Int. Conf. Comput. Vis.* pp. 962–971 (2021)
15. Jin, S., Liu, W., Xie, E., Wang, W., Qian, C., Ouyang, W., Luo, P.: Differentiable hierarchical graph grouping for multi-person pose estimation. In: *Eur. Conf. Comput. Vis.* pp. 718–734 (2020)
16. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. *Technical Report* (2009)
17. Lanchantin, J., Wang, T., Ordonez, V., Qi, Y.: General multi-label image classification with transformers. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16478–16488 (2021)

18. Li, G., Muller, M., Thabet, A., Ghanem, B.: Deepgcns: Can gcns go as deep as cnns? In: *Int. Conf. Comput. Vis.* pp. 9267–9276 (2019)
19. Li, Y., Huang, C., Loy, C.C., Tang, X.: Human attribute recognition by deep hierarchical contexts. In: *Eur. Conf. Comput. Vis.* pp. 684–700 (2016)
20. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Eur. Conf. Comput. Vis.* pp. 740–755 (2014)
21. Liu, R., Huang, J., Li, T.H., Li, G.: Causality compensated attention for contextual biased visual recognition. In: *Int. Conf. Learn. Represent.* (2022)
22. Liu, S., Zhang, L., Yang, X., Su, H., Zhu, J.: Query2label: A simple transformer way to multi-label classification. *arXiv preprint arXiv:2107.10834* (2021)
23. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *Int. Conf. Learn. Represent.* (2018)
24. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *Indian Conf. on Comput. Vis. Graph. & Image Process.* pp. 722–729. IEEE (2008)
25. Ridnik, T., Ben-Baruch, E., Zamir, N., Noy, A., Friedman, I., Protter, M., Zelnik-Manor, L.: Asymmetric loss for multi-label classification. In: *Int. Conf. Comput. Vis.* pp. 82–91 (2021)
26. Shao, J., Kang, K., Change Loy, C., Wang, X.: Deeply learned attributes for crowded scene understanding. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 4657–4666 (2015)
27. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2818–2826 (2016)
28. Wang, C., Samari, B., Kim, V.G., Chaudhuri, S., Siddiqi, K.: Affinity graph supervision for visual recognition. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8247–8255 (2020)
29. Wang, Y., He, D., Li, F., Long, X., Zhou, Z., Ma, J., Wen, S.: Multi-label classification with label graph superimposing. In: *AAAI Conf. Artif. Intell.* pp. 12265–12272 (2020)
30. Ye, J., He, J., Peng, X., Wu, W., Qiao, Y.: Attention-driven dynamic graph convolutional network for multi-label image recognition. In: *Eur. Conf. Comput. Vis.* pp. 649–665 (2020)
31. You, R., Guo, Z., Cui, L., Long, X., Bao, Y., Wen, S.: Cross-modality attention with semantic graph embedding for multi-label classification. In: *AAAI Conf. Artif. Intell.* pp. 12709–12716 (2020)
32. Zhao, J., Yan, K., Zhao, Y., Guo, X., Huang, F., Li, J.: Transformer-based dual relation graph for multi-label image recognition. In: *Int. Conf. Comput. Vis.* pp. 163–172 (2021)
33. Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., Sun, M.: Graph neural networks: A review of methods and applications. *AI Open* **1**, 57–81 (2020)
34. Zhu, F., Li, H., Ouyang, W., Yu, N., Wang, X.: Learning spatial regularization with image-level supervisions for multi-label image classification. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5513–5522 (2017)
35. Zhu, G., Zhang, L., Jiang, Y., Dang, Y., Hou, H., Shen, P., Feng, M., Zhao, X., Miao, Q., Shah, S.A.A., et al.: Scene graph generation: A comprehensive survey. *Neurocomputing* (2024)