

Panoramic Human Activity Recognition

Ruize Han¹✉, Haomin Yan^{1*}, Jiacheng Li^{1*}✉,
Songmiao Wang¹, Wei Feng^{1†}✉, and Song Wang^{2†}✉

¹ Intelligence and Computing College, Tianjin University, China

² University of South Carolina, Columbia, USA

{han_ruize, yan_hm, threeswords, smwang, wfeng}@tju.edu.cn,
songwang@cec.sc.edu

Abstract. To obtain a more comprehensive activity understanding for a crowded scene, in this paper, we propose a new problem of panoramic human activity recognition (PAR), which aims to simultaneously achieve the the recognition of individual actions, social group activities, and global activities. This is a challenging yet practical problem in real-world applications. To track this problem, we develop a novel hierarchical graph neural network to progressively represent and model the multi-granular human activities and mutual social relations for a crowd of people. We further build a benchmark to evaluate the proposed method and other related methods. Experimental results verify the rationality of the proposed PAR problem, the effectiveness of our method and the usefulness of the benchmark. We have released the source code and benchmark to the public for promoting the study on this problem.

Keywords: human action, social group, group activity, video surveillance

1 Introduction

Video-based human activity understanding is an important computer vision task, which has various practical applications in real world, e.g., video surveillance and social scene analysis [56,20]. In the past decade, this challenging task has been drawing much research interest in the computer-vision community. As shown in Fig. 1, previous works on human activity recognition can be divided into three categories. 1) Human action recognition aims to recognize the action categories of individual persons in a video [27,55,8]. 2) Human interaction recognition is proposed to recognize the human-human interactions [41]. 3) Group activity recognition is to recognize the overall activity of a group of people [23,49,25]. The last one focuses on a crowd of people while the former two commonly pay attention to the videos containing only one or a few people.

In this paper, we focus on a more comprehensive human activity understanding in the crowded scenes – we are interested in not only the overall activity of the crowd in the scene, but also the instance action of each person (referred

*Equal Contribution. †Corresponding Authors.



Fig. 1. Examples of different types of activity recognition – (a) Action recognition. (b) Interaction recognition. (c-d) Group activity recognition. (e) The proposed panoramic human activity recognition (PAR), in which the labels beside each human bounding box (green) denote the individual actions, beside each human group box (pink dotted) denote the social activities, and the label at the top left corner of the image denotes the global activity.

to as subject in this paper) and the social activity among a subset of subjects in the crowd. To achieve such a *comprehensive and multi-granular human activity understanding*, we propose a new problem – **Panoramic Human Activity Recognition (PAR)** that integrates these three sub-tasks. As shown in Fig. 1 (e), instance action recognition (Task I) aims to recognize the action of each subject in the scene, social group activity recognition (Task II) aims to divide the crowd into social groups and recognize the activity of each group, and global activity recognition (Task III) depicts the general and abstract activity according to the majority of the people. The proposed PAR problem studies the human activity understanding in multiple spatial granularity (individual, group and global) with different levels of activity categories.

The three tasks in PAR are complementary to each other and the solution of one task may benefit the others. For example, Task I recognizes the atomic action of each individual, which provides useful information for Task II of social interaction recognition, e.g., a person is talking and a person is listening may indicate the social activity of conversation if they are facing each other. Similarly, the overall activity recognition can be better achieved if we know all the individual actions and social group activities in the scene. On the contrary, the overall activity recognition provides useful priors to assist the social group activity and individual action understanding. In this paper, we aim to develop a joint framework to simultaneously address these three tasks. Compared to previous human activity understanding tasks, the proposed PAR is more challenging. A key problem is to establish a unified framework that can jointly handle all the sub tasks together, rather than address them severally or one after another. For this purpose, we expect to excavate and leverage the dependence among the different tasks and make them promote each other.

In this paper, we develop a one-stage end-to-end hierarchical graph network with a nested graph structure for PAR, in which the nodes at different hierar-

chies fitly represent the individual, group and global activities at different levels. Specifically, we first build a graph network by modeling each individual subject as a graph node. We then propose an AiO (all in one) feature aggregation module to aggregate individual feature nodes to group nodes in a bottom-to-up way. Similarly, the individual and group nodes are further aggregated into the global node. We further use a top-to-down feedback strategy for representation fusion and enhancement. With the hierarchical network architecture, we apply the multi-level supervisions for the multiple tasks in our problem.

The main contributions of this work are summarized as below:

1. We propose a new problem of Panoramic Human Activity Recognition (PAR), which aims to simultaneously recognize the individual human actions, social group activities and global activity in a crowded multi-person scene.
2. We develop a one-stage framework with hierarchical graph network that can effectively and collectively represent and model the activities at different level of granularities and mutual relations for multiple people in the scene.
3. We build a new video benchmark by adding new activity annotations and evaluation metrics to an existing dataset, for the proposed PAR. Experimental results verify the rationality of the proposed problem and the effectiveness of our method. We release the benchmark and code to the public at <https://github.com/RuizeHan/PAR>.

2 Related Work

Human action recognition and localization. Human action understanding is a fundamental task in computer vision. Early works mainly focus on the task of human action recognition, which takes a video including a human with specific action as input and aims to recognize the action category. This task can be also regarded as a video classification problem. Existing methods for action recognition can be divided into two categories, i.e., the appearance based [45,57,5,4] and the skeleton-based [44,22,38,11] methods. Recent studies started to focus on the action localization (also called action segmentation) task, including temporal action localization and spatial-temporal action localization. The former is defined to localize the temporal duration of the action in untrimmed videos and then recognize the action category [27,55]. The latter not only recognizes the action and localize its duration but also provides the spatial location (in terms of a bounding box) of the corresponding actor [8,43,48,31,26]. However, the video data used in these tasks are usually collected from the actions performed by the actors in the laboratory or from the movies/website videos, e.g., UCF [40], DALY [47], Hollywood2tubes [30], and AVA [17]. Among them, the videos commonly contains only one or very few humans, and the actors usually occupy the main part of the picture in each frame. Differently, this paper is focused on the activity understanding in the crowded scenes, which is more practical in many applications, e.g., video surveillance and social analysis, in the real world.

Human interaction recognition. Compared to the human action recognition, the study of human-human interaction recognition is less studied. Existing

human interaction recognition mainly focuses on the interactive activity involving two subjects, e.g., shaking, hugging, which occurs more in the scenarios analysis and video surveillance [35,15,52]. Similar to human action recognition, most works on interaction recognition focus on the videos collected from movies [28] and TV shows [33]. Some recent works begin to study the human interaction in the multi-person scenes. For example, in [56], a new problem of spatial-temporal human-human interaction detection is studied in a crowded scene. More comprehensive introduction to human interaction recognition can be found in a recent survey paper [41]. The classical human interaction as discussed above commonly considers the interactive activities involving two humans. Differently, in this paper we detect the social activities where the exact number of involved persons is priorly unknown – we first divide the people in the scene into human (social) groups with different sizes and then study the interactions in each group.

Social group/activity detection. Social group detection task aims to divide a crowd of people into different (sub-)groups by the social activities or relations. Early methods for this task include the group-based methods without considering each individual person [36,9], the individual-based methods aggregating the information of all individual subjects [2,39,14] and the combined methods considering both of them [32,1]. Recently, several deep learning based methods [10,6,46] are developed for the group detection task. In the recent PANDA benchmark [46], human social interaction is treated as auxiliary task for group detection in the crowded scenes. Also, a couple of recent works [6,7] aims to detect the social sub-groups in the multi-person scenes and meanwhile recognize the social activity in each sub-group. In [6,7], the social activity in each group is simply regarded as the individual action that is performed by most humans in this group – this is not practical in many real-world cases. In this paper, the task II is focused on the social group detection and its activity recognition, which is complementary to the other two tasks for more comprehensive multi-human activity understanding.

Group activity recognition. Group activity recognition (GAR) is another task for human activity understanding, which aims to recognize the activity for a group of people. Early researches directly take the video recording the activity of a group of people as input, and output the activity category [37,23], which is more like a video classification task. Recent works found that the individual human actions and the human-human interactions can help GAR. This way, several methods begin to include the individual action labels as auxiliary supervision for the GAR task [25,6]. Several other methods [23,49] propose to model the relations among multiple actors for better representation in the GAR task. This problem is actually the task III, i.e., the global activity recognition, in the proposed PAR. Difference lies in that we also simultaneously handle the other two tasks of the individual and social-group activity recognition.

Overall, the proposed PAR problem studies the human activity understanding in multiple spatial granules (individual, group and global) with different levels of activity categories, which is underexplored and has important applications.

3 The Proposed Method

3.1 Overview

In this work, we aim to jointly learn the individual actions, social group activities (based on the group detection) and the global activity in a unified one-stage framework. This way, we propose a hierarchical graph network that can well represent and model the multiple tasks. The architecture of the whole network is shown in Fig. 2.

Specifically, given a video recording a crowded scene, we extract the feature of each subject using a deep neural network. For each frame, we first model all the subjects appearing in the scene as a graph, in which each individual subject (feature) is modeled as a graph node, and each edge encodes relation between two subjects. Based on this graph, we propose a novel hierarchical graph network architecture. The proposed method can aggregate the *individual nodes* into the *group nodes* using the proposed bottom-to-up AiO (all in one) feature aggregation module. Similarly, the individual and group nodes are further aggregated into a *global node*. We also apply a top-to-down feedback strategy to further boost the mutual promotion among different tasks through the hierarchical network. We will elaborate on the proposed method in the following subsections.

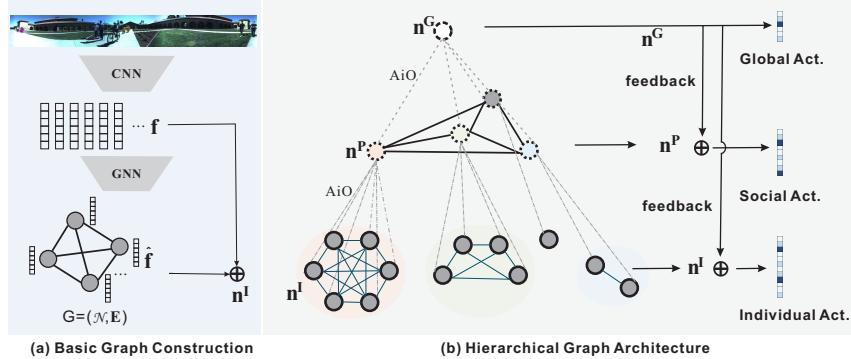


Fig. 2. Illustration of the framework of our method, which is mainly composed of the basic graph construction (as discussed in Section 3.2), and the hierarchical graph network (as discussed in Section 3.3).

3.2 Basic graph construction

We first build a fully-connected graph $G = (\mathcal{N}, \mathbf{E})$ to represent the N subjects in the scene, in which \mathcal{N} denotes the node set and \mathbf{E} encodes the relations for each pair of nodes.

Feature extraction. Similar to many previous work for action recognition [49], we employ the classical CNN network, i.e., Inception-v3 [42] followed

by the RoIAlign technique [21], to extract the deep appearance features of each subject and resize them into the same size. After that, we apply a fully connected (FC) layer to obtain the 1,024-dimension appearance feature for each subject. In the following, we denote the appearance feature of the subject u as $\mathbf{f}_u \in \mathbb{R}^{1,024}$. We use the GCN (Graph Convolutional Network) for the basic graph construction, which contains two parts, i.e., graph edge representation and graph node representation.

Graph edge representation. We use affinity matrix $\mathbf{E} \in \mathbb{R}^{N \times N}$ to learn the affinities among all the nodes in the graph, which is updated according to the node features as

$$e_{u,v} = F_1(\mathbf{f}_u) \cdot (F_2(\mathbf{f}_v))^T, u, v \in \mathcal{N}, \quad (1)$$

where \mathbf{f}_u and \mathbf{f}_v denote the feature vectors of nodes u and v , respectively. F_1 and F_2 denote the MLP (multilayer perception) networks with the same structure but unshared parameters. The output vectors are multiplied into an affinity weight $e_{u,v}$, and the affinity matrix $\mathbf{E} = [e_{u,v}]_{u,v} \in \mathbb{R}^{N \times N}$ encodes the affinity weights among all node pairs. We finally apply a softmax operation on each row of \mathbf{E} for normalization, since it is used as the weights for graph node representation updating.

Graph node representation. We then update graph node feature \mathbf{f}_u through all the connected nodes weighted by the affinity matrix as

$$\hat{\mathbf{f}}_u = F_n(\sum_v e_{u,v} \mathbf{f}_v), u, v \in \mathcal{N}, \quad (2)$$

where F_n denotes the node update network. Finally, we take the updated feature $\hat{\mathbf{f}}_u$ together with the residual connection to the original feature \mathbf{f}_u as the *individual node representation* \mathbf{n}_u^I , i.e.,

$$\mathbf{n}_u^I = \mathbf{f}_u \oplus \hat{\mathbf{f}}_u, \quad (3)$$

which embeds both the original individual feature from \mathbf{f}_u and the surrounding-aware information from $\hat{\mathbf{f}}_u$, as shown in Fig. 2(a). This way, we have finished the basic graph construction.

3.3 Hierarchical Graph Network Architecture

Bottom-to-up aggregation (B2U). After constructing the graph with each subject as a node, we propose to build a hierarchical graph network to model the different-granularity activities in the proposed problem.

Individual to group aggregation. We first consider the human social group. We aim to establish the human relation matrix $\mathbf{R} \in \mathbb{R}^{N \times N}$ to represent the human social relation among the subjects, which is calculated from two aspects, i.e., the affinity matrix \mathbf{E} in the basic graph and the distance-aware affinity matrix. Specifically, we first calculate the spatial distance matrix $\mathbf{D} \in \mathbb{R}^{N \times N}$ to encode the spatial distance between each two subject as

$$\mathbf{D}(u, v) = \frac{\sqrt{(x_u - x_v)^2 + (y_u - y_v)^2}}{\sqrt{S_u + S_v}}, \quad (4)$$

where x_u, y_u denote the coordinate of subject u (midpoint at bottom edge of the bounding box), S_u denotes the area of the bounding box of subject u . Here we take the bounding box area into consideration since the principle of near-large and far-small during imaging. We then get the relation matrix \mathbf{R} as

$$\mathbf{R} = \lambda \mathbf{E} \odot \bar{\mathbf{D}} \oplus (1 - \lambda) \check{\mathbf{D}}, \quad (5)$$

where $\check{\mathbf{D}} = \text{sigmod}(\frac{1}{\mathbf{D}})$ denotes the distance-aware affinity matrix. Similar with [49], we additionally apply a distance mask $\bar{\mathbf{D}}$ on \mathbf{E} to filter the connections between two subjects that are far from each other with a threshold ρ . The distance mask matrix $\bar{\mathbf{D}}$ is computed as

$$\bar{\mathbf{D}}(x) = \begin{cases} \mathbf{D}(x), & \text{if } \mathbf{D}(x) \leq \rho, \\ -\infty, & \text{if } \mathbf{D}(x) > \rho. \end{cases} \quad (6)$$

With the human relation matrix, we can get the human group division results through a post-processing method, e.g., a clustering algorithm. As shown in Fig. 2(b), we then aggregate the individual nodes \mathbf{n}_u^I in each group \mathcal{G}_k into the group node representation

$$\mathbf{n}_k^P = \text{AiO}(\mathbf{n}_u^I \mid u \in \mathcal{G}_k), \quad (7)$$

where $\text{AiO}(\cdot)$ denotes the proposed all in one (AiO) aggregation module, which will be discussed in detail later.

Group to global aggregation. Similar to the individual to group aggregation, as shown in Fig. 2(b), we also aggregate the individual and group nodes into a global node, as

$$\mathbf{n}^G = \text{AiO}(\mathbf{n}_u^I, \mathbf{n}_k^P \mid u \notin \forall \mathcal{G}, \forall k), \quad (8)$$

where we take both the individual node \mathbf{n}_u^I not in any group \mathcal{G} , and the aggregated group nodes \mathbf{n}_k^P in Eq. (7), as the units for aggregation. This global node can be used for representing the global activity of all subjects in the scene.

All in one (AiO) aggregation module. We then present the all in one (AiO) aggregation module used in the above. Take the individual to group aggregation using the AiO module in Eq. (7) for example. We assume the individual nodes $\mathbf{n}_u^I, u \in \mathcal{G}_k$ in a group \mathcal{G}_k are aggregated as a group node, we first build a local GCN to get the aggregation matrix $\mathbf{W} \in \mathbb{R}^{|\mathcal{G}_k| \times |\mathcal{G}_k|}$, which is similar with the graph affinity matrix in Eq. (2), and $|\mathcal{G}_k|$ denotes the number of subjects in this group. We then accumulate the values in each column of \mathbf{W} and get a weight vector $\mathbf{w} \in \mathbb{R}^{1 \times |\mathcal{G}_k|}$. We finally aggregate the individual node features with this weight vector as

$$\mathbf{n}_k^P = \sum_u (w_u \mathbf{n}_u^I \mid u \in \mathcal{G}_k), \quad (9)$$

where \mathbf{n}_k^P denotes the *aggregated node feature* of group \mathcal{G}_k , w_u is the u -th element in \mathbf{w} .

The group to global aggregation in Eq. (8) takes both the individual and group nodes as input and outputs a global node, which can be similarly achieved by the above AiO module.

Top-to-down feedback (T2D). We use the bottom-to-up hierarchical graph network to model the multi-granular activities. To further improve the mutual promotion among different tasks, we apply a top-to-down feedback strategy. Specifically, as shown in Fig. 2(b), we integrate the group node feature together with the individual/group node feature for the final representation and action category prediction as

$$\mathbf{a}_u^I = F_i(\mathbf{n}_u^I, \mathbf{n}^G), \quad \mathbf{a}_k^P = F_p(\mathbf{n}_k^P, \mathbf{n}^G), \quad \mathbf{a}^G = F_g(\mathbf{n}^G), \quad (10)$$

where \mathbf{a}_u^I , \mathbf{a}_k^P and \mathbf{a}^G are the predicted individual, social (group) and global activities, respectively, F_i , F_p , and F_g are the corresponding readout functions implemented by the MLP neural networks.

Multi-level multi-task supervisions. We use the multi-level losses as supervisions for the multiple tasks in our problem. The total loss is defined as

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_i + \mathcal{L}_p + \mathcal{L}_g + \mathcal{L}_d \\ &= \sum_u L(\mathbf{a}_u^I, \tilde{\mathbf{a}}_u^I) + \sum_k L(\mathbf{a}_k^P, \tilde{\mathbf{a}}_k^P) + L(\mathbf{a}^G, \tilde{\mathbf{a}}^G) + L(\mathbf{R}, \tilde{\mathbf{R}}), \end{aligned} \quad (11)$$

where \mathcal{L}_i , \mathcal{L}_p , \mathcal{L}_g , \mathcal{L}_d denote the losses for the individual, social group, global activity recognition, and the group detection tasks, respectively. The notations with $\tilde{\cdot}$ denote the corresponding ground-truth labels. A $\tilde{\mathbf{R}} \in \mathbb{R}^{N \times N}$ is the human group relation matrix taking the values of 0 or 1, where 1 denotes the corresponding two subjects are in the same group.

3.4 Implementation Details

Network details. The MLP networks F_1 and F_2 in Eq. (1), and F_n in Eq. (2) are all implemented by single-layer FC networks. The parameter λ in Eq. (5) is set as 0.5. The parameter ρ in Eq. (6) is set as the width of the input image with a ratio of 0.2. The readout function F_i , F_p in Eq. (10) are implemented by the three-layer FC networks, and F_g by a two-layer FC network. We use the binary cross entropy as the loss function in Eq. (11). As a new problem, in this work, we use the annotated human bounding boxes as input to alleviate the interference from the false human detection, which is common for the crowded scene in our problem. In our method, we do not integrate the temporal information of the individual and group along the video, since this need the multi-object tracking and group evolution detection results as auxiliaries. The challenging scenes make them not easy to be obtained, and the involved errors will have an impact on the main task.

Network training. During the training stage, we use the ground-truth human group division for individual to group aggregation in Eq. (7). We implement the proposed network with the PyTorch framework on the GTX 3090 GPU. The batch size is set as 4 in the experiments. We use stochastic gradient descent (SGD) algorithm with Adam method for training the network, which is trained for about 50 epochs with the learning rate of 2×10^{-5} to be convergent.

Inference. During the inference stage, we get the human group division results through a post-processing method. Specifically, with the predicted human relation matrix \mathbf{R} , we apply a self-tuning spectral clustering algorithm [53] that uses \mathbf{R} as input and automatically estimates the number of clustering groups, to obtain the group detection results.

4 Experiments

4.1 Datasets

Previous datasets for human action or group activity recognition can not meet the requirements of the proposed problem. We build a new benchmark for the proposed task. The proposed dataset is based on a state-of-the-art dataset JRDB [29] for 2D/3D person detection and tracking, which uses a mobile robot to capture the 360° RGB videos for crowded multi-person scenes, such as those at the campus, canteen, and classroom, etc. JRDB has provided the human bounding boxes with IDs as annotations. Based on it, a more recent dataset JRDB-Act [7] adds the annotations of individual human action and social group detection. The group activity label in JRDB-Act is simply defined as the combination of the involved individual action labels or the selection of them considering the occurrence frequency in the group, which seems not reasonable enough in practice. In PAR problem, the group activity is related to the individual actions but their label candidate sets are different. For example, for a two-person conversation scene, in [7] the group activity is labeled as ‘listening to someone, talking to someone’, which is the combination of the contained individual action labels. In our setting, we label such group activity as ‘chatting’, which is regarding to the whole group but not each individual. This way, we inherit the human detection annotations in JRDB and the individual action and group division annotations in JRDB-Act, and further manually annotate the social group activities and global activities for the proposed Panoramic Human Activity Recognition (PAR) task, which constitutes a new dataset – JRDB-PAR.

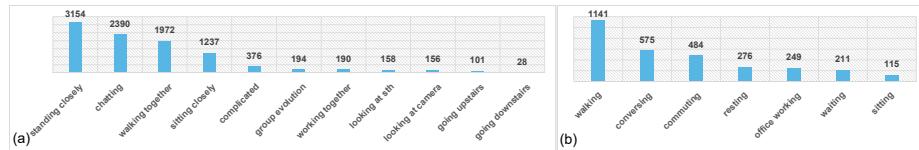


Fig. 3. Data distributions of social (a) and global activities (b) in JRDB-PAR.

JRDB-PAR contains 27 categories of individual actions, e.g., walking, talking, which is same as JRDB-Act, and 11 categories of social group activities and 7 categories of global activities. The distributions of social/global activities are shown in Fig. 3. In total, JRDB-PAR includes 27 videos, which are splitted as

20 for training and 7 for testing, following the training/validation splitting in JRDB dataset. According to statistics, JRDB-PAR contains 27,920 frames with over 628k human bounding boxes. Following the setting in [7], we select the uniformly sampled key frames (one key frame in each 15 frames) for annotation and evaluation, which is same as previous classical activity recognition datasets like CAD [3], volleyball [24]. The numbers of frames, groups and subjects in the whole dataset and key frames are summarized in Table 1. Other details for annotated labels, e.g., the number of labels for individual actions, social activities and global activities are also shown in Table 1. Note that, we adopt the multi-class labels for the activity annotation, i.e., each individual/group/frame is with one or multiple activity labels.

Table 1. Statistics results of the panoramic human activity recognition dataset.

	All Frames			Key Frames			Activity Labels		
	Frm.	Gro.	Sub.	Frm.	Gro.	Sub.	Indiv.	Social	Global
Train.	21,724	86,949	467k	1,439	5,163	27k	57,341	7,874	2,316
Test.	6,196	31,035	160k	411	1,709	9k	18,019	2,082	735
All	27,920	117,984	628k	1,850	6,872	36k	75,360	9,956	3,051

4.2 Metrics

Protocol I. To evaluate the *individual action detection*, following the previous work [16] for multi-label classification task, we adopt the commonly used metrics – precision, recall and F_1 score (denoted as \mathcal{P}_i , \mathcal{R}_i , and \mathcal{F}_i) as the evaluation metrics, which measure the individual action classification accuracy for each instance in the testing dataset.

Protocol II. *Social (group) activity recognition* includes the group detection and the activity category recognition. For the group detection, we use the classical Half metrics for performance evaluation, where the group member $\text{IoU} > 0.5$ in the predicted group and ground-truth group is taken as the true detected group, which is followed by the general protocol in group detection task [46]. Note that, we only consider the groups containing more than one subjects in this protocol. For the true detected group under the above metric, we further consider their group activity recognition results. The true detected groups with the correct activity category prediction are taken as the true social group activity predictions. Under this protocol, we calculate the precision, recall and F_1 score (denoted as \mathcal{P}_p , \mathcal{R}_p , and \mathcal{F}_p) as the social activity recognition metrics.

Protocol III. *Global activity recognition* can be also taken as a multi-label classification problem, and we apply the precision, recall, and F_1 score (denoted as \mathcal{P}_g , \mathcal{R}_g , \mathcal{F}_g) for evaluation.

The overall metric for the panoramic activity detection task is the comprehensive results of the above three metrics. Here we simply compute the average value of the above F_1 scores, i.e., $\mathcal{F}_a = \frac{1}{3}(\mathcal{F}_i + \mathcal{F}_p + \mathcal{F}_g)$ as the overall F_1 score.

4.3 Results

Comparison methods. For the proposed PAR problem, we can not find many methods for direct comparison. We try to include more state-of-the-art methods with necessary modifications as the comparison methods. ARG is a well-known method [49] for group activity recognition, which builds a learnable graph structure to model the human relation graph. SA-GAT [6] proposes to learn the human actions, sub-group and group activities together as a multi-stage multi-task problem, which uses the classical group activity recognition dataset CAD [3] with new sub-group activity annotations for evaluation. JRDB-Base [7] is the baseline method for the JRDB-Act dataset, which is similar to [6] and uses a progressive multi-loss strategy to recognize the different types of activities. Note that, the original version of ARG can not provide the group detection and social group activity recognition. We apply the group clustering algorithm used in our method and previous works [6], on the relation matrix generated by ARG for group detection. With the group detection results, we further employ the feature fusion method in [7] on ARG for the social group activity recognition task.

We also include several state-of-the-art methods for individual and group activity recognition for comparison, i.e., AT [13], SACRF [34], Dynamic [51], HiGCIN [50]. These methods can not provide the group detection and social activity recognition results, which also do not include the relation matrix in ARG. For comparison, we additionally provide the ground-truth group detection results as input to these methods and employ the feature fusion strategy in [7] to obtain the social group activity recognition results.

Table 2. Comparative results of the panoramic human activity recognition (%).

Method	Individual Act.			Social Act.			Global Act.			Overall
	\mathcal{P}_i	\mathcal{R}_i	\mathcal{F}_i	\mathcal{P}_p	\mathcal{R}_p	\mathcal{F}_p	\mathcal{P}_g	\mathcal{R}_g	\mathcal{F}_g	
ARG [49]	39.9	30.7	33.2	8.7	8.0	8.2	63.6	44.3	50.7	30.7
SA-GAT [6]	44.8	40.4	40.3	8.8	8.9	8.8	36.7	29.9	31.4	26.8
JRDB-Base [7]	19.1	34.4	23.6	14.3	12.2	12.8	44.6	46.8	45.1	27.2
Ours	51.0	40.5	43.4	24.7	26.0	24.8	52.8	31.8	38.8	35.6
AT*[13]	38.9	33.9	34.6	32.5	32.3	32.0	21.2	19.1	19.8	28.8
SACRF*[34]	31.3	23.6	25.9	25.7	24.5	24.8	42.9	35.5	37.6	29.5
Dynamic*[51]	40.7	33.4	35.1	33.5	30.1	30.9	37.5	27.1	30.6	32.2
HiGCIN*[50]	34.6	26.4	28.6	34.2	31.8	32.2	39.3	30.1	33.1	31.3
ARG*[49]	42.7	34.7	36.6	27.4	26.1	26.2	26.9	21.5	23.3	28.8
SA-GAT*[6]	39.6	34.5	35.0	32.5	32.5	30.7	28.6	24.0	25.5	30.4
JRDB-Base*[7]	21.5	44.9	27.7	54.3	45.9	48.5	38.4	33.1	34.8	37.0
Ours*	54.3	44.2	46.9	50.3	52.5	50.1	42.1	24.5	30.3	42.4

* denotes that we additionally provide the ground-truth group detection results as input.

Comparative Results. We show the performance of the proposed method and other comparative methods in Table 2. From the top half part, we can

first see that the proposed method outperforms the comparative methods on the overall score. More specifically, we can see that the proposed method achieves particularly good performance on the individual and group activity recognition. ARG performs well on the group activity recognition task, which may benefit from its flexible actor relation graphs and sparse temporal sampling strategy.

We further compare the proposed method with other state-of-the-art methods in the bottom half part of Table 2. We can first see that, with the ground-truth group detection as input, the comparison methods perform better for the social activity detection task, which outperform the proposed method. But the overall scores of most comparison methods are still lower than our method. If we also provide the ground-truth group detection in our method, i.e., ‘Ours*’, the performance will be further improved as shown in the last row.

4.4 Ablation Study

We evaluate the variations of our method by removing some components.

- w/o \mathbf{f} / $\tilde{\mathbf{f}}$: For the individual node representation, we remove the original individual feature \mathbf{f}_u or the surrounding-aware feature $\tilde{\mathbf{f}}_u$ in Eq. (3), respectively.
- w Euclid. dist.: We replace the calculation method of subject distance in Eq. (4) with the commonly used Euclidean distance.
- w/o $\check{\mathbf{D}}$ / \mathbf{E} : We remove the distance-aware matrix $\check{\mathbf{D}}$ or the edge matrix \mathbf{E} in the graph network in Eq. (5), respectively.
- w/o AiO : We replace the AiO aggregation module in Eqs. (7) and (8) with the max-pooling operation used in [7,6].
- w/o G.2I. / G.2P.: We remove the global node feature in the individual (G.2I.) or group (G.2P.) node feature representation in Eq. (10), respectively.

Table 3. Ablation study of the proposed method with its variations (%).

Method	Individual Act.			Social Act.			Global Act.			Overall
	\mathcal{P}_i	\mathcal{R}_i	\mathcal{F}_i	\mathcal{P}_p	\mathcal{R}_p	\mathcal{F}_p	\mathcal{P}_g	\mathcal{R}_g	\mathcal{F}_g	
w/o \mathbf{f}	35.6	27.2	29.6	17.4	15.9	16.4	40.2	23.8	29.3	25.1
w/o $\tilde{\mathbf{f}}$	29.3	21.7	23.9	9.8	9.2	9.4	43.1	27.0	32.4	21.9
w Euclid. dist.	43.2	32.5	35.6	12.7	12.1	12.3	41.6	24.5	30.2	26.0
w/o $\check{\mathbf{D}}$	34.6	25.8	28.2	11.3	10.8	11.0	24.8	16.1	19.0	19.4
w/o \mathbf{E}	47.8	37.8	40.6	21.0	20.9	20.6	33.6	19.5	24.2	28.4
w/o AiO	19.8	14.6	16.1	18.8	16.7	17.3	34.2	30.4	31.3	21.6
w/o G.2I.	31.1	22.9	25.3	12.5	10.9	11.4	75.4	45.0	55.2	30.6
w/o G.2P.	36.4	27.5	30.0	11.9	10.4	10.9	78.4	46.8	57.3	32.7
Ours	51.0	40.5	43.4	24.7	26.0	24.8	52.8	31.8	38.8	35.6

We show the ablation study results in Table 3. We can first see that original individual feature \mathbf{f}_u and the surrounding node embedded feature $\tilde{\mathbf{f}}_u$ are both

important for individual node representation. Especially the latter $\hat{\mathbf{f}}_u$ is very useful to the social activity recognition.

Next, in the bottom-to-up (B2U) aggregation stage, we first investigate the effect of the spatial-aware matrix $\check{\mathbf{D}}$ and the edge matrix \mathbf{E} in relation matrix learning. We can see that $\check{\mathbf{D}}$ is vital in relation modeling. It is easy to explain that the social relations among the people in the crowd are highly related to the spatial distance. But the simple Euclidean distance between bounding box center performs not well enough. Also, only using the spatial-aware matrix, i.e., w/o \mathbf{E} , can not get the performance as using both of them. We then investigate the effectiveness of the AiO module in the proposed B2U aggregation, we can see that the aggregation method AiO is verified to be effective in our method.

Finally, in the top-to-down (T2D) feedback stage, we can see that the individual action recognition results get worse when removing the global node feature in individual node representation. Similarly, the social activity recognition performance also gets worse without the global node feature embedding. This demonstrates the effectiveness of the T2D feedback strategy in our method.

4.5 Experimental Analysis

Individual/global activity recognition. We evaluate the tasks of individual action and global activity recognition, where we follow the original settings in the comparative GAR methods without modification. We show the results in Table 4. We can see that, although the margin is not very large, the proposed method performs better than the comparative methods on these two tasks. We can also find that the performance on JRDB-PAR is generally low, where is great potential for more methods to be developed.

Table 4. Comparisons of the individual action and global activity recognition (%).

Method	Individual Act.			Global Act.		
	\mathcal{P}_i	\mathcal{R}_i	\mathcal{F}_i	\mathcal{P}_g	\mathcal{R}_g	\mathcal{F}_g
AT [13]	36.8	30.1	31.7	17.4	15.7	16.1
SACRF [34]	39.2	29.4	32.2	34.8	26.2	28.4
Dynamic [51]	46.6	37.7	39.7	31.9	23.7	26.4
HiGCIN [50]	36.9	30.1	31.6	46.0	34.2	38.0
Ours	51.0	40.5	43.4	52.8	31.8	38.8

Group detection. We further solely evaluate the group detection results in Table 5. We first use the classical Half metric used in group detection task [46], i.e., the predicted group with the group member $\text{IoU} > 0.5$ with the ground-truth group is taken as the true detected group, which is denoted as $\text{IOU}@0.5$. We extend this metric by increasing the threshold 0.5 to 1 with a step of 0.1, which means improving the criteria bar. We then plot the accuracy curve using each threshold and the corresponding accuracy. This way, we can compute the

AUC (Area Under the Curve) score as the metric namely IOU@AUC. With the predicted and ground-truth binary human group relation matrices $\bar{\mathbf{R}}$ and $\tilde{\mathbf{R}}$, where 1 denotes the corresponding two subjects are in the same group, we compute the matrix IOU score as $\text{Mat.IOU} = \frac{\sum \text{AND}(\mathbf{R}, \tilde{\mathbf{R}})}{\sum \text{OR}(\mathbf{R}, \tilde{\mathbf{R}})}$, where AND and OR denotes the functions of element-wise logical and/or operations.

As shown in Table 5, we first apply two straightforward approaches for comparison. Here ‘Dis.Mat + LP’ denotes that we use the distance matrix among the subjects followed by a label propagation method [54] algorithm as the post-processing method to get the group relation matrix. ‘GNN + GRU’ uses a GNN to model the group relations among the subjects followed by a classical GRU model to integrate the temporal information. We also include other methods, i.e., ARG [49], SA-GAT [6], and JRDG-Base [7], that can provide the group detection results for comparison. We can see that, the proposed method outperforms the comparison methods with a large margin under all metrics.

Table 5. Comparative group detection results (%).

	IOU@0.5	IOU@AUC	Mat.IOU
Dis.Mat + LP [54]	33.4	14.1	12.9
GNN + GRU	34.5	21.7	20.1
ARG [49]	35.2	21.6	19.3
SA-GAT [6]	29.1	20.4	16.6
JRDG-Base [7]	38.4	26.3	20.6
Ours	53.9	38.1	30.6

5 Conclusion

In this paper, we have studied a new problem – panoramic human activity recognition (PAR), which is more comprehensive, challenging and practical than previous action recognition tasks in crowded scene video analysis. To handle this problem, we propose a novel hierarchical neural network composed of the basic graph construction, bottom-to-up node aggregation, and top-to-down feedback strategy, to model the multi-granular activities. Based on an existing dataset, we build a benchmark for this new problem. Experimental results are very promising and verify the rationality of this new problem and the effectiveness of the proposed baseline method. In the future, we plan to integrate the temporal information, e.g., the multiple object tracking [19,18] and group evolution detection, into the proposed method. An interesting problem maybe to investigate whether we can simultaneously handle the multiple object tracking [12] and PAR and make them help with each other.

Acknowledgment. This work was supported by the National Natural Science Foundation of China under Grants U1803264, 62072334, and the Tianjin Research Innovation Project for Postgraduate Students under Grant 2021YJSB174.

References

1. Bazzani, L., Cristani, M., Murino, V.: Decentralized particle filter for joint individual-group tracking. In: CVPR (2012)
2. Chang, M.C., Krahnstoever, N., Ge, W.: Probabilistic group-level motion analysis and scenario recognition. In: ICCV (2011)
3. Choi W, Shahid K, S.S.: What are they doing?: Collective activity classification using spatio-temporal relationship among people. In: ICCV (2009)
4. Diba, A., Fayyaz, M., Sharma, V., Arzani, M.M., Yousefzadeh, R., Gall, J., Van Gool, L.: Spatio-temporal channel correlation networks for action classification. In: ECCV (2018)
5. Du, Y., Yuan, C., Li, B., Zhao, L., Li, Y., Hu, W.: Interaction-aware spatio-temporal pyramid attention networks for action classification. In: ECCV (2018)
6. Ehsanpour, M., Abedin, A., Saleh, F., Shi, J., Reid, I., Rezatofighi, H.: Joint learning of social groups, individuals action and sub-group activities in videos. In: ECCV (2020)
7. Ehsanpour, M., Saleh, F.S., Savarese, S., Reid, I.D., Rezatofighi, H.: Jrdb-act: A large-scale dataset for spatio-temporal action, social group and activity detection. In: arXiv preprint (2021)
8. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: ICCV (2019)
9. Feldmann, M., Fränken, D., Koch, W.: Tracking of extended objects and group targets using random matrices. IEEE Transactions on Signal Processing **59**(4), 1409–1420 (2010)
10. Fernando, T., Denman, S., Sridharan, S., Fookes, C.: Gd-gan: Generative adversarial networks for trajectory prediction and group detection in crowds. In: ACCV (2018)
11. Friji, R., Drira, H., Chaieb, F., Kchok, H., Kurtek, S.: Geometric deep neural network using rigid and non-rigid transformations for human action recognition. In: ICCV (2021)
12. Gan, Y., Han, R., Yin, L., Feng, W., Wang, S.: Self-supervised multi-view multi-human association and tracking. In: ACM MM (2021)
13. Gavrillyuk, K., Sanford, R., Javan, M., Snoek, C.G.: Actor-transformers for group activity recognition. In: CVPR (2020)
14. Ge, W., Collins, R.T., Ruback, R.B.: Vision-based analysis of small groups in pedestrian crowds. IEEE TPAMI **34**(5), 1003–1016 (2012)
15. Gemeren, C.V., Poppe, R., Veltkamp, R.C.: Spatio-temporal detection of fine-grained dyadic human interactions. In: International Workshop on Human Behavior Understanding (2016)
16. Godbole, S., Sarawagi, S.: Discriminative methods for multi-labeled classification. In: Pacific-Asia conference on knowledge discovery and data mining (2004)
17. Gu, C., Sun, C., Ross, D.A., Vondrick, C., Pantofaru, C., Li, Y., Vijayanarasimhan, S., Toderici, G., Ricco, S., Sukthankar, R., et al.: Ava: A video dataset of spatio-temporally localized atomic visual actions. In: CVPR (2018)
18. Han, R., Feng, W., Zhang, Y., Zhao, J., Wang, S.: Multiple human association and tracking from egocentric and complementary top views. IEEE TPAMI (2021). <https://doi.org/10.1109/TPAMI.2021.3070562>
19. Han, R., Feng, W., Zhao, J., Niu, Z., Zhang, Y., Wan, L., Wang, S.: Complementary-view multiple human tracking. In: AAAI (2020)

20. Han, R., Zhao, J., Feng, W., Gan, Y., Wan, L., Wang, S.: Complementary-view co-interest person detection. In: ACM MM (2020)
21. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV (2017)
22. Huang, Z., Wan, C., Probst, T., Van Gool, L.: Deep learning on lie groups for skeleton-based action recognition. In: CVPR (2017)
23. Ibrahim, M.S., Mori, G.: Hierarchical relational networks for group activity recognition and retrieval. In: ECCV (2018)
24. Ibrahim, M.S., Muralidharan, S., Deng, Z., Vahdat, A., Mori, G.: A hierarchical deep temporal model for group activity recognition. In: CVPR (2016)
25. Ibrahim M S, M.G.: Hierarchical relational networks for group activity recognition and retrieval. In: ECCV (2018)
26. Li, Y., Chen, L., He, R., Wang, Z., Wu, G., Wang, L.: Multisports: A multi-person video dataset of spatio-temporally localized sports actions. In: arXiv preprint (2021)
27. Ma, F., Zhu, L., Yang, Y., Zha, S., Kundu, G., Feiszli, M., Shou, Z.: Sf-net: Single-frame supervision for temporal action localization. In: ECCV (2020)
28. Marszalek, M., Laptev, I., Schmid, C.: Actions in context. In: CVPR (2009)
29. Martin-Martin, R., Patel, M., Rezatofighi, H., Shenoi, A., Gwak, J., Frankel, E., Sadeghian, A., Savarese, S.: Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments. IEEE TPAMI (2021). <https://doi.org/10.1109/TPAMI.2021.3070543>
30. Mettes, Van Gemert J C, S.C.G.M.: Spot on: Action localization from pointily-supervised proposals. In: ECCV (2016)
31. Pan, J., Chen, S., Shou, M.Z., Liu, Y., Shao, J., Li, H.: Actor-context-actor relation network for spatio-temporal action localization. In: CVPR (2021)
32. Pang, S.K., Li, J., Godsill, S.J.: Detection and tracking of coordinated groups. IEEE Transactions on Aerospace and Electronic Systems **47**(1), 472–502 (2011)
33. Patron-Perez, A., Marszalek, M., Reid, I., Zisserman, A.: Structured learning of human interactions in tv shows. IEEE TPAMI **34**(12), 2441–2453 (2012)
34. Pramono, R.R.A., Chen, Y.T., Fang, W.H.: Empowering relational network by self-attention augmented conditional random fields for group activity recognition. In: ECCV (2020)
35. Ryoo, M.S., Aggarwal, J.K.: Spatio-temporal relationship match: Video structure comparison for recognition of complex human activities. In: ICCV (2009)
36. Shao, J., Change Loy, C., Wang, X.: Scene-independent group profiling in crowd. In: CVPR (2014)
37. Shu, T., Todorovic, S., Zhu, S.C.: Cern: Confidence-energy recurrent network for group activity recognition. In: CVPR (2017)
38. Si, C., Jing, Y., Wang, W., Wang, L., Tan, T.: Skeleton-based action recognition with spatial reasoning and temporal stack learning. In: ECCV (2018)
39. Solera, F., Calderara, S., Cucchiara, R.: Socially constrained structural learning for groups detection in crowd. IEEE TPAMI **38**(5), 995–1008 (2015)
40. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. Computer Science (2012)
41. Stergiou, A., Poppe, R.: Analyzing human-human interactions: A survey. Computer vision and image understanding **188**(Nov.), 102799.1–102799.12 (2019)
42. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016)
43. Tang, J., Xia, J., Mu, X., Pang, B., Lu, C.: Asynchronous interaction aggregation for action detection. In: ECCV (2020)

44. Vemulapalli, R., Arrate, F., Chellappa, R.: Human action recognition by representing 3d skeletons as points in a lie group. In: CVPR (2014)
45. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Gool, L.V.: Temporal segment networks: Towards good practices for deep action recognition. In: ECCV (2016)
46. Wang, X., Zhang, X., Zhu, Y., Guo, Y., Yuan, X., Xiang, L., Wang, Z., Ding, G., Brady, D., Dai, Q., et al.: Panda: A gigapixel-level human-centric video dataset. In: CVPR (2020)
47. Weinzaepfel, P., Martin, X., Schmid, C.: Towards weakly-supervised action localization. In: arXiv preprint (2016)
48. Wu, J., Kuang, Z., Wang, L., Zhang, W., Wu, G.: Context-aware rcnn: A baseline for action detection in videos. In: ECCV (2020)
49. Wu, J., Wang, L., Wang, L., Guo, J., Wu, G.: Learning actor relation graphs for group activity recognition. In: CVPR (2019)
50. Yan, R., Xie, L., Tang, J., Shu, X., Tian, Q.: Higcin: Hierarchical graph-based cross inference network for group activity recognition. IEEE TPAMI (2020). <https://doi.org/10.1109/TPAMI.2020.3034233>
51. Yuan, H., Ni, D.: Learning visual context for group activity recognition. In: AAAI (2021)
52. Yun, K., Honorio, J., Chattopadhyay, D., Berg, T.L., Samaras, D.: Two-person interaction detection using body-pose features and multiple instance learning. In: CVPRW (2012)
53. Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: NeurIPS (2004)
54. Zhan, X., Liu, Z., Yan, J., Lin, D., Loy, C.C.: Consensus-driven propagation in massive unlabeled data for face recognition. In: ECCV (2018)
55. Zhang, X.Y., Shi, H., Li, C., Li, P.: Multi-instance multi-label action recognition and localization based on spatio-temporal pre-trimming for untrimmed videos. In: AAAI (2020)
56. Zhao, J., Han, R., Gan, Y., Wan, L., Feng, W., Wang, S.: Human identification and interaction detection in cross-view multi-person videos with wearable cameras. In: ACM MM (2020)
57. Zhou, Y., Sun, X., Zha, Z.J., Zeng, W.: Mict: Mixed 3d/2d convolutional tube for human action recognition. In: CVPR (2018)