

Hunting Group Clues with Transformers for Social Group Activity Recognition

Masato Tamura[✉], Rahul Vishwakarma[✉], and Ravigopal Vennelakanti

Hitachi America, Ltd.

`masato.tamura.sf@hitachi.com`,

`{rahul.vishwakarma,ravigopal.vennelakanti}@hal.hitachi.com`

Abstract. This paper presents a novel framework for social group activity recognition. As an expanded task of group activity recognition, social group activity recognition requires recognizing multiple sub-group activities and identifying group members. Most existing methods tackle both tasks by refining region features and then summarizing them into activity features. Such heuristic feature design renders the effectiveness of features susceptible to incomplete person localization and disregards the importance of scene contexts. Furthermore, region features are sub-optimal to identify group members because the features may be dominated by those of people in the regions and have different semantics. To overcome these drawbacks, we propose to leverage attention modules in transformers to generate effective social group features. Our method is designed in such a way that the attention modules identify and then aggregate features relevant to social group activities, generating an effective feature for each social group. Group member information is embedded into the features and thus accessed by feed-forward networks. The outputs of feed-forward networks represent groups so concisely that group members can be identified with simple Hungarian matching between groups and individuals. Experimental results show that our method outperforms state-of-the-art methods on the Volleyball and Collective Activity datasets.

Keywords: social group activity recognition, group activity recognition, social scene understanding, attention mechanism, transformer

1 Introduction

Social group activity recognition is a task of recognizing multiple sub-group activities and identifying group members in a scene. This task is derived from group activity recognition, which needs to recognize only one group activity in a scene. Both tasks have gained tremendous attention in recent years for potential applications such as sports video analysis, crowd behavior analysis, and social scene understanding [1–5, 12–14, 17–19, 22, 24–28, 34, 35, 37, 40, 43–50]. In the context of these tasks, the term “action” denotes an atomic movement of a single person, and the term “activity” refers to a more complex relation of movements performed by a group of people. Although our framework can recognize both actions and activities, we focus on group activities.

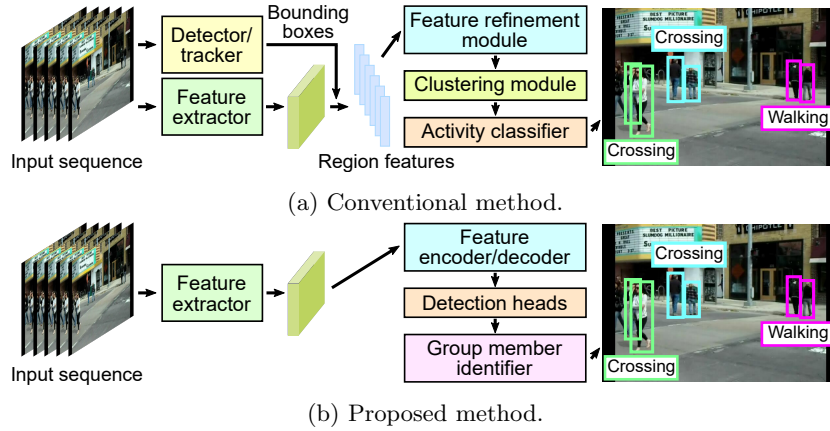


Fig. 1: Overviews of conventional and proposed social group activity recognition methods. The labels in the right image show predicted social group activities.

Most existing methods decompose the recognition process into two independent parts; person localization and activity recognition (See Fig. 1a) [5, 12–14, 17, 19, 22, 27, 34, 35, 37, 43, 44, 49, 50]. Person localization identifies regions where people are observed in a scene with bounding boxes. These boxes are used to extract region features from feature maps. The region features are further refined to encode spatio-temporal relations with refinement modules such as recurrent neural networks (RNNs) [8, 16], graph neural networks (GNNs) [21, 42], and transformers [41]. The refined features are summarized for activity recognition.

While these methods have demonstrated significant improvement, they have several drawbacks attributed to the heuristic nature of feature design. Since region features are extracted from bounding box regions in feature maps, the effectiveness of the features is affected by the localization performance. Most existing methods ignore this effect and evaluate their performances with region features of ground truth boxes. However, several works [5, 13, 35, 44] show that the recognition performance is slightly degraded when using predicted boxes instead of ground truth boxes. Moreover, substantial scene contexts are discarded by using region features because they are typically dominated by features of the people in the boxes. Scene contexts such as object positions and background situations are sometimes crucial to recognize group activities. For instance, the positions of sports balls are informative to recognize group activities in sports games. These features should be leveraged to enhance recognition performance.

Another challenge specific to social group activity recognition is that utilizing region features is sub-optimal to identify group members. Ehsanpour *et al.* [13] use region features as node features of graph attention networks (GATs) [42] and train them to output adjacency matrices that have low probabilities for people in different groups and high probabilities for those in the same groups. During inference, spectral clustering [32] is applied to the adjacency matrices to divide

people into groups. Because adjacency matrices reflect semantic similarities of node features, this method may not work if region features of people in the same group have different semantics such as doing different actions.

To address these challenges, we propose a novel social group activity recognition method that can be applied to both social group activity recognition and group activity recognition. We leverage a transformer-based object detection framework [6, 52] to obviate the need for the heuristic feature design in existing methods (See Fig. 1b). Attention modules in transformers play crucial roles in our method. We design our method in such a way that the attention modules identify and then aggregate features relevant to social group activities, generating an effective feature for each social group. Because activity and group member information is embedded into the generated features, the information can be accessed by feed-forward networks (FFNs) in the detection heads. The outputs of the detection heads are designed so concisely that group member identification can be performed with simple Hungarian matching between groups and individuals. This identification method differs from Ehsanpour *et al.*'s method [13] in that their method relies on individuals' features to divide people into groups, while our method generates features that are embedded with clues for grouping people, enabling effective group identification.

To summarize, our contributions are three-fold: (1) We propose a novel social group activity recognition method that leverages the attention modules to generate effective social group features. (2) Our method achieves better or competitive performance to state-of-the-art methods on both group activity recognition and social group activity recognition. (3) We perform comprehensive analyses to reveal how our method works with activities under various conditions.

2 Related Works

2.1 Group Member Identification and Activity Recognition

Several group member identification methods [15, 33, 38] were proposed independently of group activity recognition. They utilize hand-crafted features to find interactions and identify group members on the basis of their interactions.

In group activity recognition, deep-neural-network-based methods have become dominant due to the learning capability of the networks. Ibrahim *et al.* [19] proposed an RNN-based method that uses convolutional neural networks to extract region features and long short-term memories to refine the features. This architecture captures the temporal dynamics of each person between frames and the spatial dynamics of people in a scene. After their work, several RNN-based methods were proposed [5, 22, 35, 37, 43].

GNNs are also utilized to model the spatio-temporal context and relationships of people in a scene. Wu *et al.* [44] used graph convolutional networks (GCNs) [21] to capture spatio-temporal relations of people's appearances and positions between frames. Ehsanpour *et al.* [13] adopted GATs [42] to learn underlying interactions and divide people into social groups with adjacency matrices. Hu *et al.* [17] utilized both RNNs and GNNs with reinforcement learning

to refine features. Yuan *et al.* [49] used person-specific dynamic graphs that dynamically change connections of GNNs for each node.

Several works introduced transformers into group activity recognition. Gavrielyuk *et al.* [14] used transformer encoders to refine region features. Li *et al.* [27] proposed spatial-temporal transformers that can encode spatio-temporal dependence and decode the group activity information. Zhou *et al.* [50] proposed multi-scale spatio-temporal stacked transformers for compositional understanding and relational reasoning in group activities.

Our method differs from existing methods in that they rely on region features, while our method generates social group features with the attention modules in transformers, resulting in improving the performance.

2.2 Detection Transformer

Carion *et al.* [6] proposed a transformer-based object detector called DETR, which regards object detection as a set prediction problem. One significant difference between conventional object detectors and DETR is that conventional ones need heuristic detection points whose features are used to predict object classes and bounding boxes, while DETR obviates such heuristic components by letting queries in transformer decoders aggregate features for their target objects. DETR shows competitive performance compared with conventional state-of-the-art detectors even without such heuristic components.

To further improve the performance of DETR, several methods have been proposed [11, 39, 52]. Zhu *et al.* [52] proposed Deformable DETR that replaces standard transformers with deformable ones. Deformable attention modules in the transformers combine a sparse sampling of deformable convolution [10] and dynamic weighting of standard attention modules, which significantly reduces the computational complexity of the attention weight calculation. This reduction allows Deformable DETR to use multi-scale feature maps from backbone networks. To leverage non-heuristic designs and multi-scale feature maps, we use deformable transformers to generate social group features.

3 Proposed Method

We leverage a deformable-transformer-based object detection framework [52], whose details are omitted due to the limited space. We encourage readers to refer to the paper for more details.

3.1 Overall Architecture

Figure 2 shows the overall architecture of the proposed method. Given a frame sequence $\mathbf{x} \in \mathbb{R}^{3 \times T \times H \times W}$, a feature extractor extracts a set of multi-scale feature maps $\mathbf{Z}_f = \{\mathbf{z}_i^{(f)} \mid \mathbf{z}_i^{(f)} \in \mathbb{R}^{D_i \times T \times H'_i \times W'_i}\}_{i=1}^{L_f}$, where T is the length of the sequence, H and W are the height and width of the frame, H'_i and W'_i are those of the output feature maps, D_i is the number of channels, and L_f is the number

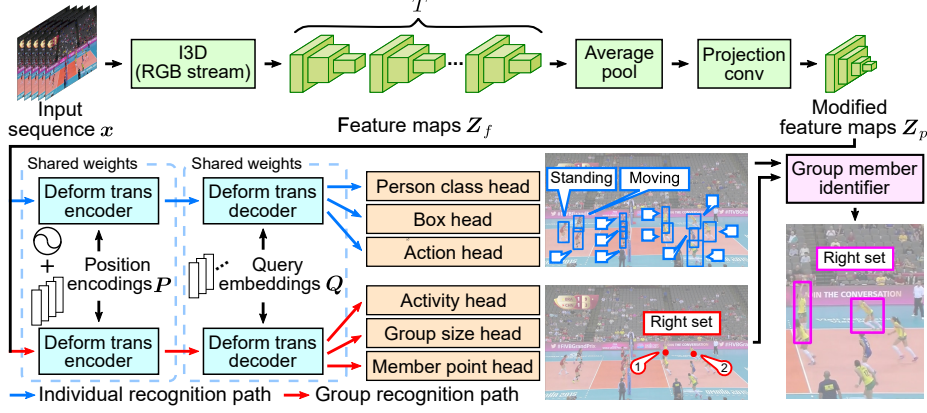


Fig. 2: Overall architecture of the proposed method.

of scales. We adopt the inflated 3D (I3D) network [7] as a feature extractor to embed local spatio-temporal context into feature maps. Note that we use only the RGB stream of I3D because group members are identified by their positions, which cannot be predicted with the optical flow stream. To reduce the computational costs of transformers, each feature map $z_i^{(f)}$ is mean-pooled over the temporal dimension and input to a projection convolution layer that reduces the channel dimension from D_i to D_p . One additional projection convolution layer with a kernel size of 3×3 and stride of 2×2 is applied to the smallest feature map to further add the scale.

Features in the modified feature maps are refined and aggregated with deformable transformers. Given a set of the modified multi-scale feature maps $Z_p = \{z_i^{(p)} \mid z_i^{(p)} \in \mathbb{R}^{D_p \times H'_i \times W'_i}\}_{i=1}^{L_f+1}$, a set of refined feature maps $Z_e = \{z_i^{(e)} \mid z_i^{(e)} \in \mathbb{R}^{D_p \times H'_i \times W'_i}\}_{i=1}^{L_f+1}$ is obtained as $Z_e = f_{enc}(Z_p, P)$, where $f_{enc}(\cdot, \cdot)$ is stacked deformable transformer encoder layers and $P = \{p_i \mid p_i \in \mathbb{R}^{D_p \times H'_i \times W'_i}\}_{i=1}^{L_f+1}$ is a set of multi-scale position encodings [52], which supplement the attention modules with position and scale information to identify where each feature lies in the feature maps. The encoder helps features to acquire rich social group context by exchanging information in a feature map and between multi-scale feature maps. These enriched feature maps are fed into the deformable transformer decoder to aggregate features. Given a set of refined feature maps Z_e and learnable query embeddings $Q = \{q_i \mid q_i \in \mathbb{R}^{2D_p}\}_{i=1}^{N_q}$, a set of feature embeddings $H = \{h_i \mid h_i \in \mathbb{R}^{D_p}\}_{i=1}^{N_q}$ is obtained as $H = f_{dec}(Z_e, Q)$, where N_q is the number of query embeddings and $f_{dec}(\cdot, \cdot)$ is stacked deformable transformer decoder layers. Each decoder layer predicts locations that contain features relevant to input embeddings and aggregates the features from the locations with the dynamic weighting. We design queries in such a way that one query captures at most one social group. This design enables each query to aggregate features of its target social group from the refined feature maps.

The feature embeddings are transformed into prediction results with detection heads. Here we denote the localization results in normalized image coordinates. Social group activities are recognized by predicting activities and identifying group members. The identification is performed with a group size head and group member point head. The size head predicts the number of people in a target social group, and the point head indicates group members by localizing the centers of group members' bounding boxes. This design enables our method to identify group members with simple point matching during inference as described in Sec. 3.3. The predictions of activity class probabilities $\{\hat{\mathbf{v}}_i \mid \hat{\mathbf{v}}_i \in [0, 1]^{N_v}\}_{i=1}^{N_q}$, group sizes $\{\hat{s}_i \mid \hat{s}_i \in [0, 1]\}_{i=1}^{N_q}$, and sequences of group member points $\{\hat{\mathbf{U}}_i\}_{i=1}^{N_q}$ are obtained as $\hat{\mathbf{v}}_i = f_v(\mathbf{h}_i)$, $\hat{s}_i = f_s(\mathbf{h}_i)$, and $\hat{\mathbf{U}}_i = f_u(\mathbf{h}_i, \mathbf{r}_i)$, where N_v is the number of activity classes, $\hat{\mathbf{U}}_i = (\hat{\mathbf{u}}_j^{(i)} \mid \hat{\mathbf{u}}_j^{(i)} \in [0, 1]^2)_{j=1}^M$ is a sequence of points that indicate centers of group members' bounding boxes, M is a hyper-parameter that defines the maximum group size, $f_v(\cdot)$, $f_s(\cdot)$, and $f_u(\cdot, \cdot)$ are the detection heads for each prediction, and $\mathbf{r}_i \in [0, 1]^2$ is a reference point, which is used in the same way as the localization in Deformable DETR [52]. The predicted group sizes are values normalized with M . All the detection heads are composed of FFNs with subsequent sigmoid functions. We describe the details of the detection heads in the supplementary material.

Individual recognition can be performed by replacing the group recognition heads with individual recognition heads. We empirically find that using different parameters of deformable transformers for individual recognition and social group recognition does not show performance improvement. Therefore, we use shared parameters to reduce computational costs. The details of the individual recognition heads are described in the supplementary material.

3.2 Loss Calculation

We view social group activity recognition as a direct set prediction problem and match predictions and ground truths with the Hungarian algorithm [23] during training following the training procedure of DETR [6]. The optimal assignment is determined by calculating the matching cost with the predicted activity class probabilities, group sizes, and group member points. Given a ground truth set of social group activity recognition, the set is first padded with $\phi^{(gr)}$ (no activity) to change the set size to N_q . With the padded ground truth set, the matching cost of the i -th element in the ground truth set and j -th element in the prediction set is calculated as follows:

$$\mathcal{H}_{i,j}^{(gr)} = \mathbb{1}_{\{i \notin \Phi^{(gr)}\}} \left[\eta_v \mathcal{H}_{i,j}^{(v)} + \eta_s \mathcal{H}_{i,j}^{(s)} + \eta_u \mathcal{H}_{i,j}^{(u)} \right], \quad (1)$$

$$\mathcal{H}_{i,j}^{(v)} = - \frac{\mathbf{v}_i^T \hat{\mathbf{v}}_j + (\mathbf{1} - \mathbf{v}_i)^T (\mathbf{1} - \hat{\mathbf{v}}_j)}{N_v}, \quad (2)$$

$$\mathcal{H}_{i,j}^{(s)} = |s_i - \hat{s}_j|, \quad (3)$$

$$\mathcal{H}_{i,j}^{(u)} = \frac{\sum_{k=1}^{S_i} \|\mathbf{u}_k^{(i)} - \hat{\mathbf{u}}_k^{(j)}\|_1}{S_i}, \quad (4)$$

where $\Phi^{(gr)}$ is a set of ground-truth indices that correspond to $\phi^{(gr)}$, $\mathbf{v}_i \in \{0, 1\}^{N_v}$ is a ground truth activity label, $s_i \in [0, 1]$ is a ground truth group size normalized with M , S_i is an unnormalized ground truth group size, $\mathbf{u}_k^{(i)} \in [0, 1]^2$ is a ground truth group member point normalized with the image size, and $\eta_{\{v,s,u\}}$ are hyper-parameters. Group member points in the sequence $\mathbf{U}_i = (\mathbf{u}_k^{(i)})_{k=1}^{S_i}$ are sorted in ascending order along X coordinates as seen from the image of the group recognition result in Fig. 2. We use this arrangement because group members are typically seen side by side at the same vertical positions in an image, and the order of group member points is clear from their positions, which makes the prediction easy. We evaluate the performances with other arrangements and compare the results in Sec. 4.4. Using Hungarian algorithm, the optimal assignment is calculated as $\hat{\omega}^{(gr)} = \arg \min_{\omega \in \Omega_{N_q}} \sum_{i=1}^{N_q} \mathcal{H}_{i,\omega(i)}^{(gr)}$, where Ω_{N_q} is the set of all possible permutations of N_q elements.

The training loss for social group activity recognition \mathcal{L}_{gr} is calculated between matched ground truths and predictions as follows:

$$\mathcal{L}_{gr} = \lambda_v \mathcal{L}_v + \lambda_s \mathcal{L}_s + \lambda_u \mathcal{L}_u, \quad (5)$$

$$\mathcal{L}_v = \frac{1}{|\Phi^{(gr)}|} \sum_{i=1}^{N_q} [\mathbb{1}_{\{i \notin \Phi^{(gr)}\}} l_f(\mathbf{v}_i, \hat{\mathbf{v}}_{\hat{\omega}^{(gr)}(i)}) + \mathbb{1}_{\{i \in \Phi^{(gr)}\}} l_f(\mathbf{0}, \hat{\mathbf{v}}_{\hat{\omega}^{(gr)}(i)})], \quad (6)$$

$$\mathcal{L}_s = \frac{1}{|\Phi^{(gr)}|} \sum_{i=1}^{N_q} \mathbb{1}_{\{i \notin \Phi^{(gr)}\}} |s_i - \hat{s}_{\hat{\omega}^{(gr)}(i)}|, \quad (7)$$

$$\mathcal{L}_u = \frac{1}{|\Phi^{(gr)}|} \sum_{i=1}^{N_q} \sum_{j=1}^{S_i} \mathbb{1}_{\{i \notin \Phi^{(gr)}\}} \left\| \mathbf{u}_j^{(i)} - \hat{\mathbf{u}}_j^{(\hat{\omega}^{(gr)}(i))} \right\|_1, \quad (8)$$

where $\lambda_{\{v,s,u\}}$ are hyper-parameters and $l_f(\cdot, \cdot)$ is the element-wise focal loss function [29] whose hyper-parameters are described in [51].

Individual recognition is jointly learned by matching ground truths and predictions of person class probabilities, bounding boxes, and action class probabilities and calculating the losses between matched ground truths and predictions. The matching and loss calculations are performed by slightly modifying the original matching costs and losses of Deformable DETR [52]. We describe the details of these matching and loss calculations in the supplementary material.

3.3 Group Member Identification

The outputs of the detection heads represent groups in group sizes and group member points that indicate the centers of group members' bounding boxes. These values have to be transformed into values that indicate individuals. We transform the predicted values into indices that refer to the elements in the individual prediction set with the following simple process during inference. To match the group member points and individual predictions, the Hungarian algorithm [23] is used instead of just calculating the closest center of a bounding box

for each group member point. Hungarian algorithm can prevent multiple group member points from matching the same individuals and thus slightly improve the performance. The matching cost between i -th group member point of k -th social group prediction and j -th individual prediction is calculated as follows:

$$\mathcal{H}_{i,j}^{(gm,k)} = \frac{\|\hat{\mathbf{u}}_i^{(k)} - f_{cent}(\hat{\mathbf{b}}_j)\|_2}{\hat{c}_j}, \quad (9)$$

where $\hat{\mathbf{b}}_j \in [0, 1]^4$ is a predicted bounding box of an individual, $\hat{c}_j \in [0, 1]$ is a detection score of the individual, and $f_{cent}(\cdot)$ is a function that calculates the center of a bounding box. By applying the Hungarian algorithm to this matching cost, the optimal assignment is calculated as $\hat{\omega}^{(gm,k)} = \arg \min_{\omega \in \Omega_{N_q}} \sum_{i=1}^{\lfloor M \times \hat{s}_k \rfloor} \mathcal{H}_{i,\omega(i)}^{(gm,k)}$, where $\lfloor \cdot \rfloor$ rounds an input value to the nearest integer. Finally, the index set of individuals for k -th social group prediction is obtained as $\mathbf{G}_k = \{\hat{\omega}^{(gm,k)}(i)\}_{i=1}^{\lfloor M \times \hat{s}_k \rfloor}$.

4 Experiments

4.1 Datasets and Evaluation Metrics

We evaluate the performance of our method on two publicly available benchmark datasets: the Volleyball dataset [19] and Collective Activity dataset [9]. The Volleyball dataset contains 4,830 videos of 55 volleyball matches, which are split into 3,493 training videos and 1,337 test videos. The center frame of each video is annotated with bounding boxes, actions, and one group activity. The number of action and activity classes are 9 and 8, respectively. Because the original annotations do not contain group member information, we use an extra annotation set provided by Sendo and Ukita [36]. We combine the original annotations with the group annotations in the extra set and use them for our experiments. Note that annotations other than the group annotations in the extra set are not used for a fair comparison. The Collective Activity dataset contains 44 videos of life scenes, which are split into 32 training videos and 12 test videos. The videos are annotated every ten frames with bounding boxes and actions. The group activity is defined as the action with the largest number in the scenes. The number of action classes is 6. Because the original annotations do not have group member information, Ehsanpour *et al.* [13] annotated group labels. We use their annotations for our experiments.

We divide the evaluation into two parts: group activity recognition and social group activity recognition. In the evaluation of group activity recognition, we follow the detection-based settings [5, 13, 35, 44] and use classification accuracy as an evaluation metric. Because our method is designed to predict multiple group activities, we need to select one from them for group activity recognition. We choose the predicted activity of the highest probability and compare it with the ground truth activity. In the evaluation of social group activity recognition, different metrics are used for each dataset because each scene in the Volleyball

dataset contains only one social group activity, while that in the Collective Activity dataset contains multiple social group activities. For the Volleyball dataset, group identification accuracy is used as an evaluation metric. One group prediction is first selected in the same way as group activity recognition, and then the predicted bounding boxes of the group members are compared with the ground truth boxes. The selected prediction results are correct if the predicted activity is correct and the predicted boxes have IoUs larger than 0.5 with the corresponding ground truth boxes. For the Collective Activity dataset, mAP is used as an evaluation metric. Prediction results are judged as true positives if the predicted activities are correct, and all the predicted boxes of the group members have IoUs larger than 0.5 with the corresponding ground truth boxes.

4.2 Implementation Details

We use the RGB stream of I3D [7] as a backbone feature extractor and input features from *Mixed_3c*, *Mixed_4f*, and *Mixed_5c* layers into the deformable transformers. The hyper-parameters of the deformable transformers are set in accordance with the setting of Deformable DETR [52], where $L_f = 3$, $D_p = 256$, and $N_q = 300$. We initialize I3D with the parameters trained on the Kinetics dataset [20] and deformable transformers with the parameters trained on the COCO dataset [30]. We use the AdamW [31] optimizer with the batch size of 16, the initial learning rate of 10^{-4} , and the weight decay of 10^{-4} . Training epochs are set to 120, and the learning rate is decayed after 100 epochs. We set the length of the sequence T to 9. Ground truth labels of the center frame are used to calculate the losses. To augment the training data, we randomly shift frames in the temporal direction and use bounding boxes from visual trackers as ground truth boxes when a non-annotated frame is at the center. We also augment the training data by random horizontal flipping, scaling, and cropping. Following the DETR’s training [6], auxiliary losses are used to boost the performance. The maximum group size M is set to 12. The hyper-parameters are set as $\eta_v = \lambda_v = 2$, $\eta_s = \lambda_s = 1$, and $\eta_u = \lambda_u = 5$.

While evaluating performances with the Collective Activity dataset, some specific settings are used. For the evaluation of group activity recognition, training epochs are set to 10, and the learning rate is decayed after 5 epochs because the losses converge in a few epochs due to the limited diversity of the scenes in the dataset. For the evaluation of social group activity recognition, the length of the sequence T is set to 17 following the setting of Ehsanpour *et al.* [13].

4.3 Group Activity Recognition

Comparison against State-of-the-Art. We compare our method against state-of-the-art methods on group activity recognition. Table 1 shows the comparison results. The values without the brackets demonstrate the detection-based performances, while those inside the brackets indicate the performances with ground truth bounding boxes. We show the performances of individual action recognition for future reference. Several detection-based performances are not

Table 1: Comparison of group activity recognition. The values w/ and w/o the brackets show the performances in the ground-truth-based and detection-based settings, respectively. Full comparison results are in the supplementary material.

Method	Volleyball				Collective Activity			
	Activity		Action		Activity		Action	
SSU [5]	86.2	(90.6)	–	(81.8)	–	(–)	–	(–)
stagNet [35]	87.6	(89.3)	–	(–)	87.9	(89.1)	–	(–)
ARG [44]	91.5	(92.5)	39.8	(83.0)	86.1	(88.1)	49.6	(77.3)
CRM [4]	–	(93.0)	–	(–)	–	(85.8)	–	(–)
Actor-Transformers [14]	–	(94.4)	–	(85.9)	–	(92.8)	–	(–)
Ehsanpour <i>et al.</i> [13]	93.0	(93.1)	41.8	(83.3)	89.4	(89.4)	55.9	(78.3)
Pramono <i>et al.</i> [34]	–	(95.0)	–	(83.1)	–	(95.2)	–	(–)
P ² CTDM [45]	–	(92.7)	–	(–)	–	(96.1)	–	(–)
DIN [49]	–	(93.6)	–	(–)	–	(95.9)	–	(–)
GroupFormer [27]	95.0 [*]	(95.7)	–	(85.6)	85.2 [*]	(87.5 [†] /96.3)	–	(–)
Ours	96.0	(–)	65.0	(–)	96.5	(–)	64.9	(–)

^{*} We evaluated the performance with the publicly available source codes.

[†] We evaluated but were not able to reproduce the reported accuracy because the configuration file for the Collective Activity dataset is not publicly available.

reported because existing works typically use ground-truth boxes for the evaluation. To compare the effectiveness of our method with these methods, we evaluate GroupFormer [27], which is the strongest baseline of group activity recognition, with predicted boxes of Deformable DETR [52]. Note that Deformable DETR is fine-tuned on each dataset for a fair comparison, which demonstrates 90.8 and 90.2 mAP on the Volleyball and Collective Activity datasets, respectively.

As seen from the table, our method outperforms state-of-the-art methods in the detection-based setting. We confirm that GroupFormer shows the performance degradation as well as the previous methods [5, 13, 35, 44] when predicted bounding boxes are used. These results indicate that the latest region-feature-based method still suffers from incomplete person localization and that our feature generation has advantages over these methods. Even compared to the ground-truth-based performances, our method shows the best performance. It is worth noting that our method uses only RGB images as inputs, while GroupFormer utilizes optical flows and poses in addition to RGB data. These results suggest that our features are more effective than region features and that it is not optimal to restrict regions of features to bounding boxes.

Analysis of Group Annotations. As described in Sec. 4.1, we use the additional group annotations to fully leverage our social group activity recognition capability. We analyze the effect of the group annotations on group activity recognition by investigating the performances of both GroupFormer [27] and our method with and without the group annotations. Note that hereinafter we use

Table 2: Analysis of the effect of the group annotations with Volleyball. The values w/ and w/o the brackets demonstrate the performances in the ground-truth-based and detection-based settings, respectively.

Method	Annotation type	Activity
GroupFormer [27]	Original	95.0 [*] (95.7)
	Group	93.2 [‡] (96.1 [*])
Ours	Original	95.0 (-)
	Group	96.0 (-)

^{*} We evaluated the performance with the publicly available source codes.

[‡] We trained a group member detector and evaluated the performance with publicly available source codes.

the Volleyball dataset for analyses because the diversity of the scenes in the Collective Activity dataset is limited. To evaluate GroupFormer with the group annotations in the detection-based setting, we trained Deformable DETR [52] with bounding boxes of only group members, which is intended to detect only people involved in activities. The detector shows the performance of 87.1 mAP. Among all the results, GroupFormer with the group annotations in the ground-truth-based setting demonstrates the best performance. However, the performance is substantially degraded when the predicted boxes are used. This is probably because group member detection underperforms and degrades the recognition performance. As our method does not rely on bounding boxes to predict group activities, the performance does not degrade even if group members cannot be identified correctly. Accordingly, our method demonstrates the best performance in the detection-based setting.

4.4 Social Group Activity Recognition

Comparison against State-of-the-Art. To demonstrate the effectiveness, we compare our method against Ehsanpour *et al.*'s method [13], which is a state-of-the-art method that tackles social group activity recognition, and GroupFormer [27], which is the strongest baseline on group activity recognition. Due to the unavailability of both Ehsanpour *et al.*'s source codes and their performance report on the Volleyball dataset, we implemented their algorithm based on our best understanding and evaluated the performance on the dataset. To detect group members in the evaluation of GroupFormer, we trained Deformable DETR [52] as described in the group annotation analysis section. Because this group member detection cannot be applied to multi-group detection, we evaluate GroupFormer only on the Volleyball dataset.

Table 3 shows the results of the Volleyball dataset. As shown in the table, our method yields significant performance gains over the other methods, which demonstrates the improvement in group member identification as well as in activity recognition. Our method aggregates features that are embedded with

Table 3: Comparison of social group activity recognition with Volleyball.

Method	Accuracy	Right				Left			
		Set	Spike	Pass	Winpoint	Set	Spike	Pass	Winpoint
Ehsanpour <i>et al.</i> [13] [§]	44.5	17.2	74.0	49.0	29.9	19.7	79.6	25.0	28.4
GroupFormer [27] [‡]	48.8	25.0	56.6	59.0	51.7	31.5	55.3	58.8	51.0
Ours	60.6	35.9	68.2	81.9	50.6	50.6	53.6	74.3	56.9

[§] Because the source codes are not publicly available, we implemented their algorithm based on our best understanding and evaluated the performance.

[‡] We trained a group member detector and evaluated the performance with publicly available source codes.

Table 4: Comparison of social group activity recognition with Collective Act.

Method	mAP	Crossing	Waiting	Queueing	Walking	Talking
Ehsanpour <i>et al.</i> [13]	51.3	–	–	–	–	–
Ours	46.0	49.2	64.5	54.1	55.6	6.56

clues for grouping people from feature maps. It is highly likely that this feature aggregation contributes to the high accuracy of identifying activities with different distributions of group members in an image. We qualitatively analyze how features are aggregated depending on the distribution of group members and discuss the analysis results towards the end of the qualitative analysis section.

The comparison results on the Collective Activity dataset are listed in Table 4. As seen from the table, Ehsanpour *et al.*’s method shows better performance than our method. We find that our method demonstrates relatively low performance on the activity “Talking”. This low performance is probably attributed to the number of samples in the training data. In the test data, 86 % of the samples with the activity “Talking” have group sizes of four, while the training data has only 57 samples whose group sizes are four, which is 0.8 % of the training data. As our method learns to predict group sizes, the number of samples in training data for each group size affects the performance. We analyze this effect in the subsequent section.

Analysis of Group Sizes. The group size prediction is one of the key factors to identify group members and thus affects social group activity recognition performance. To analyze this effect, we evaluate the performance of each group size and compare the results with Ehsanpour *et al.*’s method [13] and GroupFormer [27]. Table 5 shows the results. As shown in the table, the performances of our method are moderately correlated to the training data ratios, while the other two methods do not show the correlation. This is the drawback of our method that relies on group-size learning. However, our method shows competitive performances in both small and large group sizes if there are a certain amount of training data.

Table 5: Analysis of group sizes with Volleyball.

Method	Group size (Training data ratio)					
	1 (36 %)	2 (21 %)	3 (19 %)	4 (6 %)	5 (5 %)	6 (12 %)
Ehsanpour <i>et al.</i> [13] [§]	45.3	48.2	61.2	27.3	15.8	32.5
GroupFormer [27] [‡]	57.3	29.6	58.4	28.4	44.7	54.4
Ours	83.6	42.9	52.4	26.1	39.5	63.8

[§] Because the source codes are not publicly available, we implemented their algorithm based on our best understanding and evaluated the performance.

[‡] We trained a group member detector and evaluated the performance with publicly available source codes.

Table 6: Analysis of the order of member points with Volleyball.

Order of the group member points	Probability of changes in order	Accuracy
Ascending order in X coordinates	7.4 %	60.6
Ascending order in Y coordinates	13 %	55.5

In contrast, each of the other two methods shows competitive performances only in either large or small group sizes. These results imply that our method does not have the performance dependence on group sizes and thus can achieve high performance with large-scale training data.

Analysis of Group Member Point Order. As described in Sec. 3.2, group member points in a ground truth point sequence are sorted in ascending order along X coordinates. To confirm the effectiveness of this arrangement, we compare the performances with two arrangements. Table 6 shows the comparison results. As shown in the table, our method demonstrates better performance when group member points are sorted in ascending order along X coordinates than in ascending order along Y coordinates. The probabilities in the table indicate the ratio of the changes in the point order when small perturbations are added to ground-truth bounding box positions. The higher probability implies that the order of group member points changes more frequently when group members move. These results suggest that the order changes more frequently when group member points are sorted in ascending order along Y coordinates and that the order is difficult to predict with slight differences in box positions.

Qualitative Analysis. The deformable attention modules are the critical components to aggregate features relevant to social group activity recognition and generate social group features. To analyze how the attention modules aggregate features for various social group activities, we visualize the attention locations of the transformer decoder in Fig. 3. We show locations with the top four attention weights in the last layer of the decoder. The purple bounding boxes show the




Fig. 3: Visualization of the attention locations in the deformable transformer decoder. We show the locations of the top four attention weights. The large circles mean that the locations are in the low-resolution feature maps.

group members, the red circles show the predicted group member points, and the yellow circles show the attention locations. The small and large yellow circles mean that the locations are in the high and low-resolution feature maps, respectively, showing a rough range of image areas affecting the generated features. The figure shows that features are typically aggregated from low-resolution feature maps if group members are located in broad areas, and vice versa. These results indicate that the attention modules can effectively aggregate features depending on the distribution of group members and contribute to improving the performance of social group activity recognition.

5 Conclusions

We propose a novel social group activity recognition method that leverages deformable transformers to generate effective social group features. This feature generation obviates the need for region features and hence makes the effectiveness of the social group features person-localization-agnostic. Furthermore, the group member information extracted from the features is represented so concisely that our method can identify group members with simple Hungarian matching, resulting in high-performance social group activity recognition. We perform extensive experiments and show significant improvement over existing methods.

Acknowledgement

Computational resource of AI Bridging Cloud Infrastructure (ABCI) provided by National Institute of Advanced Industrial Science and Technology was used.

References

1. Amer, M.R., Lei, P., Todorovic, S.: HiRF: Hierarchical random field for collective activity recognition in videos. In: ECCV (September 2014)
2. Amer, M.R., Todorovic, S.: Sum product networks for activity recognition. IEEE TPAMI **38**(4), 800–813 (April 2016)
3. Amer, M.R., Todorovic, S., Fern, A., Zhu, S.C.: Monte carlo tree search for scheduling activity recognition. In: ICCV (December 2013)
4. Azar, S.M., Atigh, M.G., Nickabadi, A., Alahi, A.: Convolutional relational machine for group activity recognition. In: CVPR (June 2019)
5. Bagautdinov, T.M., Alahi, A., Fleuret, F., Fua, P.V., Savarese, S.: Social scene understanding: End-to-end multi-person action localization and collective activity recognition. In: CVPR (July 2017)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (November 2020)
7. Carreira, J., Zisserman, A.: Quo vadis, action recognition? A new model and the kinetics dataset. In: CVPR (July 2017)
8. Cho, K., van Merriënboer, B., Çaglar Gülçehre, Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. In: EMNLP (October 2014)
9. Choi, W., Shahid, K., Savarese, S.: What are they doing? : Collective activity classification using spatio-temporal relationship among people. In: ICCVW (September 2009)
10. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: ICCV (October 2017)
11. Dai, X., Chen, Y., Yang, J., Zhang, P., Yuan, L., Zhang, L.: Dynamic DETR: End-to-end object detection with dynamic attention. In: ICCV (October 2021)
12. Deng, Z., Vahdat, A., Hu, H., Mori, G.: Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In: CVPR (June 2016)
13. 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 (August 2020)
14. Gavriluk, K., Sanford, R., Javan, M., Snoek, C.G.M.: Actor-transformers for group activity recognition. In: CVPR (June 2020)
15. Ge, W., Collins, R.T., Ruback, R.B.: Vision-based analysis of small groups in pedestrian crowds. IEEE TPAMI **34**(5), 1003–1016 (2012)
16. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Computation **9**(8), 1735–1780 (November 1997)
17. Hu, G., Cui, B., He, Y., Yu, S.: Progressive relation learning for group activity recognition. In: CVPR (June 2020)
18. Ibrahim, M.S., Mori, G.: Hierarchical relational networks for group activity recognition and retrieval. In: ECCV (September 2018)
19. Ibrahim, M.S., Muralidharan, S., Deng, Z., Vahdat, A., Mori, G.: A hierarchical deep temporal model for group activity recognition. In: CVPR (June 2016)
20. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A.: The kinetics human action video dataset (May 2017), arXiv:1705.06950
21. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (April 2017)

22. Kong, L., Qin, J., Huang, D., Wang, Y., Gool, L.V.: Hierarchical attention and context modeling for group activity recognition. In: ICASSP (April 2018)
23. Kuhn, H.W., Yaw, B.: The hungarian method for the assignment problem. *Naval Res. Logist. Quart* pp. 83–97 (March 1955)
24. Lan, T., Sigal, L., Mori, G.: Social roles in hierarchical models for human activity recognition. In: CVPR (June 2012)
25. Lan, T., Wang, Y., Yang, W., Mori, G.: Beyond actions: Discriminative models for contextual group activities. In: NIPS (December 2010)
26. Lan, T., Wang, Y., Yang, W., Robinovitch, S.N., Mori, G.: Discriminative latent models for recognizing contextual group activities. *IEEE TPAMI* **34**(8), 1549–1562 (August 2012)
27. Li, S., Cao, Q., Liu, L., Yang, K., Liu, S., Hou, J., Yi, S.: GroupFormer: Group activity recognition with clustered spatial-temporal transformer. In: ICCV (October 2021)
28. Li, X., Chuah, M.C.: SBGAR: Semantics based group activity recognition. In: ICCV (October 2017)
29. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (October 2017)
30. 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: ECCV (September 2014)
31. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (May 2019)
32. Ng, A., Jordan, M., Weiss, Y.: On spectral clustering: Analysis and an algorithm. In: NIPS (December 2002)
33. Park, H., Shi, J.: Social saliency prediction. In: CVPR (June 2015)
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 (August 2020)
35. Qi, M., Qin, J., Li, A., Wang, Y., Luo, J., Gool, L.V.: stagNet: An attentive semantic rnn for group activity recognition. In: ECCV (September 2018)
36. Sendo, K., Ukita, N.: Heatmapping of people involved in group activities. In: MVA (May 2019)
37. Shu, T., Todorovic, S., Zhu, S.C.: CERN: Confidence-energy recurrent network for group activity recognition. In: CVPR (July 2017)
38. Shu, T., Xie, D., Rothrock, B., Todorovic, S., Zhu, S.: Joint inference of groups, events and human roles in aerial videos. In: CVPR (June 2015)
39. Sun, Z., Cao, S., Yang, Y., Kitani, K.M.: Rethinking transformer-based set prediction for object detection. In: ICCV (October 2021)
40. Tang, J., Shu, X., Yan, R., Zhang, L.: Coherence constrained graph lstm for group activity recognition. *IEEE TPAMI* **44**(2), 636–647 (2022)
41. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NIPS (December 2017)
42. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: ICLR (April 2018)
43. Wang, M., Ni, B., Yang, X.: Recurrent modeling of interaction context for collective activity recognition. In: CVPR (July 2017)
44. Wu, J., Wang, L., Wang, L., Guo, J., Wu, G.: Learning actor relation graphs for group activity recognition. In: CVPR (June 2019)

45. Yan, R., Shu, X., Yuan, C., Tian, Q., Tang, J.: Position-aware participation-contributed temporal dynamic model for group activity recognition. IEEE TNNLS (June 2021)
46. Yan, R., Tang, J., Shu, X., Li, Z., Tian, Q.: Participation-contributed temporal dynamic model for group activity recognition. In: ACMMM (October 2018)
47. Yan, R., Xie, L., Tang, J., Shu, X., Tian, Q.: HiGCIN: Hierarchical graph-based cross inference network for group activity recognition. IEEE TPAMI (October 2020)
48. Yan, R., Xie, L., Tang, J., Shu, X., Tian, Q.: Social adaptive module for weakly-supervised group activity recognition. In: ECCV (August 2020)
49. Yuan, H., Ni, D., Wang, M.: Spatio-temporal dynamic inference network for group activity recognition. In: ICCV (October 2021)
50. Zhou, H., Kadav, A., Shamsian, A., Geng, S., Lai, F., Zhao, L., Liu, T., Kapadia, M., Graf, H.P.: COMPOSER: Compositional learning of group activity in videos (December 2021), arXiv:2112.05892
51. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points (April 2019), arXiv:1904.07850
52. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable transformers for end-to-end object detection. In: ICLR (May 2021)