

Spatio-Temporal Dynamic Inference Network for Group Activity Recognition

Hangjie Yuan¹ Dong Ni^{1,2*} Mang Wang³

¹College of Control Science and Engineering, Zhejiang University, Hangzhou, China

²State Key Laboratory of Industrial Control Technology, Zhejiang University, Hangzhou, China

³DAMO Academy, Alibaba Group, China

{hj.yuan,dni}@zju.edu.cn, wangmang.wm@alibaba-inc.com

Abstract

Group activity recognition aims to understand the activity performed by a group of people. In order to solve it, modeling complex spatio-temporal interactions is the key. Previous methods are limited in reasoning on a predefined graph, which ignores the inherent person-specific interaction context. Moreover, they adopt inference schemes that are computationally expensive and easily result in the over-smoothing problem. In this paper, we manage to achieve spatio-temporal person-specific inferences by proposing Dynamic Inference Network (DIN), which composes of Dynamic Relation (DR) module and Dynamic Walk (DW) module. We firstly propose to initialize interaction fields on a primary spatio-temporal graph. Within each interaction field, we apply DR to predict the relation matrix and DW to predict the dynamic walk offsets in a joint-processing manner, thus forming a person-specific interaction graph. By updating features on the specific graph, a person can possess a global-level interaction field with a local initialization. Experiments indicate both modules' effectiveness. Moreover, DIN¹ achieves significant improvement compared to previous state-of-the-art methods on two popular datasets under the same setting, while costing much less computation overhead of the reasoning module.

1. Introduction

Group activity recognition (GAR) aims to infer an overall activity performed by a group of people in the scene [9, 22, 4, 45, 32, 43, 15, 48]. It has aroused research interests due to various applications, including surveillance/sports video analysis, social scene understanding, etc. The critical problem that lies in GAR is to infer a group-level activity representation given a video clip, which asks for elaborately designed reasoning modules.

*Corresponding author.

¹Codes are available at https://github.com/JacobYuan7/DIN_GAR.




Figure 1. **Examples of right set and right pass group activity.** The red bounding box annotated with a star is the person performing the key action for the activity. The grey arrow denotes the key interaction linking the starred person and the semantically important person, which is always not aligned in the spatial or temporal domain. The person indices do not start from 1 because we only illustrate part of the images.

Recently proposed reasoning modules mainly incorporate spatio-temporal interactive factors to get a refined activity representation. Modeling of agents' interactions has been widely studied. The mostly adopted methods are recurrent neural networks [1, 49], the attention mechanism [41, 19] and graph neural networks (GNNs) [38, 13, 44]. GNNs have been a frequently adopted method in GAR [32, 43, 46, 31], which performs message passing on a constructed semantic graph and achieves competitive results on publicly available benchmarks.

However, previous methods using GNNs stick to a paradigm that models the interaction between individuals on a predefined graph as shown in Figure 2. It is a feasible way but bears several drawbacks: **i)** Those who interact with a given person should be person-specific but not predefined. Like in Figure 1, a person will interact with people depending on their own context: the 8th person in the left video interacts with the 9th person who is about to spike the ball; the 10th person in the right video interacts with the 9th person who is about to set the ball. A predefined graph can not suit every person's inference. **ii)** Previous predefined graph models infer interactions on a fully-connected [43] or criss-cross [46, 31] graph which is shown in Figure




Figure 2. **Visualizations of three inference schemes** in spatio-temporal domain with GNNs for GAR. The green node denotes the feature to be updated. The purple nodes denote features involved in updating the green node. (a) Fully-connected graph inference. (b) Criss-cross graph inference. (c) Proposed person-specific dynamic graph inference, which is unique for every green node. The dashed box is an example of an initialized interaction field.

2(a) and (b). It easily results in the over-smoothing [27] that makes features indistinguishable and damages the performance. Also, it costs overmuch computation overhead if expanding to long video clips or expanding to a scenario with too many people in the scene.

Aiming at solving the drawbacks mentioned above, inspired by [11, 50], we present Dynamic Inference Network (DIN), which contains Dynamic Relation (DR) and Dynamic Walk (DW). These two modules combined can predict a person-specific interaction graph for better modeling interactions as shown in Figure 2(c). For a given person feature on a spatio-temporal graph, we set a spatio-temporal interaction field around it as an initialization, which is shared between DR and DW. This interaction field determines the people to be involved in inferring the interaction graph. The initialized field size will not increase if the spatial or temporal axis expands, which reduces the computation.

Within this initialized interaction field, we use DR to predict a person-specific relation matrix, denoting the interaction relations between persons. The features in the interaction field endow the relations with an interaction context. Then, to facilitate the model to learn from complex spatio-temporal interactions, we use DW to predict dynamic walk offsets for every feature within the field. The dynamic walks allow for the locally initialized interaction field to form a graph that enables global-level interactions. The proposed modules are easy for deployment onto any widely used backbones to form a pipeline named DIN. Besides, previous methods seldom make computational complexity analysis, which is a significant evaluation for a designed module. In this paper, we present computational complexity analysis and show that our modules cost less computation overhead while performing better.

To summarize, our contributions are listed as follows:

- We propose DIN to construct person-specific interaction graphs in the spatio-temporal domain, which are not predefined and can also serve as a general approach for modeling interactions.
- We propose DR that predicts person-specific relation

matrices and DW that allows for the locally initialized interaction field to update features globally. Both are proved useful by experiments.

- We prove by experiments that a small size of initialized interaction field is sufficient for existing datasets. We use a case visualization to exemplify that interaction graphs can capture the key person and key interactions, and a locally initialized interaction field can cover a global-level interaction field with proposed modules.
- DIN achieves state-of-the-art performances under the setting of the same backbone and input modality on two widely used benchmarks, while costing much less computation overhead of the reasoning module.

2. Related Work

Group Activity Recognition Group activity recognition was firstly proposed in [9]. Following works [10, 25, 24, 8, 2] were basically to extract hand-crafted features (e.g., HOG [12]) and apply graphical models to infer group activity representations. With the boom of deep learning, methods incorporating convolution neural networks (CNNs) and recurrent neural networks (RNNs) have proved effective. For example, the works of [22, 4] managed to model the temporal dynamics in action level or group level via RNNs on CNN features. The works of [42, 45, 32, 36] applied RNNs to the modeling of person interactions. The attention mechanism also proved its effectiveness in GAR. The works of [45, 32, 39] combined RNNs with attention mechanisms to capture the key features in the spatial or temporal domain. Specifically, the self-attention mechanism was introduced to learn the temporal evolution and spatial interactions [15, 31].

GNNs, which inferred on graph-structured data, attracted researchers' attention in GAR. ARG [43] firstly proposed to use graph convolution networks (GCNs) to learn person interactions on a spatio-temporal graph. Later, several works [46, 31, 47] improved the previous fully-connected graph to a criss-cross one when modeling relations and aggregating features. However, they all ignored the person-specific interaction context. Our work is partly inspired by deformable convolution [11, 51], whose relations are not conditioning on the person features. Moreover, related work like DGMN [50] which mentioned 'dynamic' constrained in implicit pixel-level spatial feature enhancement, while our pipeline suits video processing and agent-level spatio-temporal reasoning.

Modeling of Interactions The modeling of interactions is significant in understanding a complex system with multiple objects/agents [5, 6]. Many research areas inherently involve the modeling of interactions like trajectory prediction [16, 38, 34], human object interaction [33, 29, 14] and scene graph generation [44, 49, 7]. In GAR, modeling interactions is crucial in understanding their overall activity [4]. Among

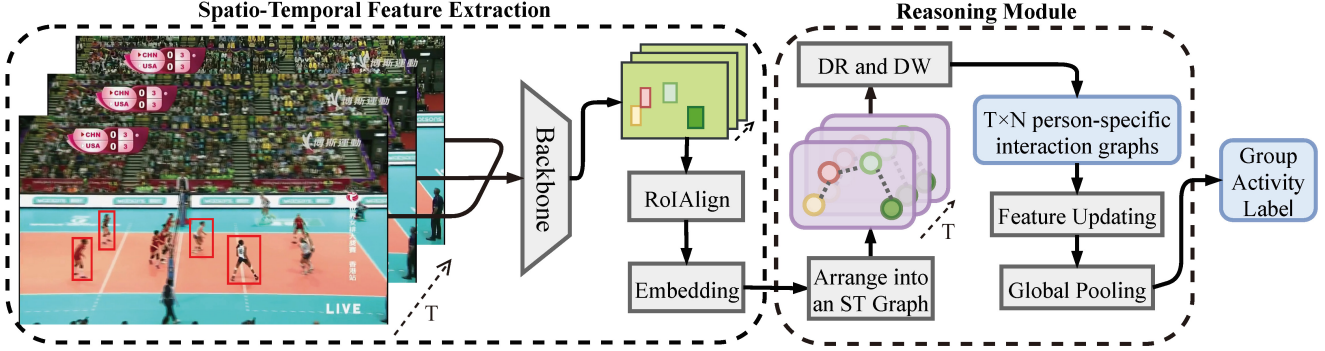


Figure 3. **The overall pipeline of Dynamic Inference Network.** Generally, it consists of two stages: i) Spatio-temporal feature extraction, ii) Reasoning module. Note that there will be $T \times N$ unique interaction graphs for updating. In our codebase, the first stage is shared with previous methods. The main variations are in the Reasoning Module. We only illustrate 4 bounding boxes in the image for clarity.

their adopted methods, GNNs have been a frequently chosen method. Some related works like EvolveGCN [30] explored a better representation learning strategy on evolving graphs and EvolveGraph [26] explored a prediction method to adjust the structure of one graph. However, we focus on the exploration of constructing dynamic agent-specific graphs based on their interaction field. The proposed modules are general approaches to tackle the modeling of interactions in related problems.

3. Method

In this section, we firstly outline the pipeline of DIN. Then, we give a brief review of previous GNN reasoning modules for GAR. Finally we introduce the modules that we propose to dynamically infer the group activity. To better present the idea, we specifically present the feature updating method for the i th person feature.

3.1. Dynamic Inference Network

The integrated framework, which we denote as Dynamic Inference Network (DIN), is illustrated in Figure 3. The DIN takes in a short clip of videos, which is fed into a selected backbone network to extract visual features. For the backbone network, we mainly experiment on ResNet-18 [18] and VGG-16 [37] to demonstrate the effectiveness of our proposed module and to seek for a fair comparison with previous methods. Then RoIAlign [17] is applied to extract the person features aligned with bounding boxes, which are then embedded to a D -dimension space. We stack the person features to form $\mathbf{X} \in \mathbb{R}^{T \times N \times D}$, where T, N denotes the temporal steps (*i.e.*, temporal dimension) and number of annotated people in each frame (*i.e.*, spatial dimension) respectively. Note that the spatial dimension is ordered by people’s coordinates following [43, 15]. It is then arranged into a Spatio-Temporal graph (ST graph). The proposed DR and DW dynamically predict a specific interaction graph for a selected feature ($T \times N$ interaction graphs in total). Thus, we can operate feature updating accordingly.

After the above inference, we can perform a global pooling to get the final group representation, which contains a max-pooling layer along the spatial dimension and an average pooling layer along the temporal dimension. The training objective is the cross-entropy loss for group activities. Although many previous methods like [32, 43, 3, 15] use additional cross-entropy loss for individual actions, the action labels are actually ill-defined [47] and expensive in labeling. We use cheap group activity labels while still achieving competitive results.

Although computational complexity analysis in Section 4 already indicates that DR and DW bring in limited parameters and FLOPs apart from the backbone and the embedding layer, we take a step further to seek for a lighter reasoning module. In practice, we apply pointwise convolution [20] before the reasoning module to reduce the dimension of \mathbf{X} from D to D_l . We name this model **Lite DIN**.

3.2. Recap of Previous GNN Reasoning

We start the recap by introducing Actor Relation Graph (ARG) in [43]. The spatio-temporal feature extraction stage for ARG is identical to the DIN as illustrated in Figure 3. It uses a fully-connected graph as illustrated in Figure 2(a). The spatial and temporal dimension of person features \mathbf{X} are collapsed to one, denoted as $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^{TN}$ with $\mathbf{x}_i \in \mathbb{R}^D$. Their pairwise relations can be denoted as $\mathbf{R} = \{r_{i,j} | i, j = 1, \dots, TN\}$ with $r_{i,j} \in \mathbb{R}^1$, which can be computed by

$$r_{i,j} = \frac{\theta(\mathbf{x}_i)^T \phi(\mathbf{x}_j)}{\sqrt{D_r}} \quad (1)$$

$$\tilde{r}_{i,j} = \text{softmax}_j(r_{i,j}) = \frac{\exp(r_{i,j})}{\sum_{q=1}^{TN} \exp(r_{i,q})} \quad (2)$$

where θ and ϕ are linear transformations functions, *i.e.*, $\theta(\mathbf{x}_i) = \mathbf{W}_\theta \mathbf{x}_i$ with $\mathbf{W}_\theta \in \mathbb{R}^{D_r \times D}$ and $\phi(\mathbf{x}_i)$ is similarly defined; D_r is the dimension of the embedding space; softmax_j defines a softmax function along the index j to get the normalized relations $\tilde{r}_{i,j}$. We do not formulate distance mask here for clarity.

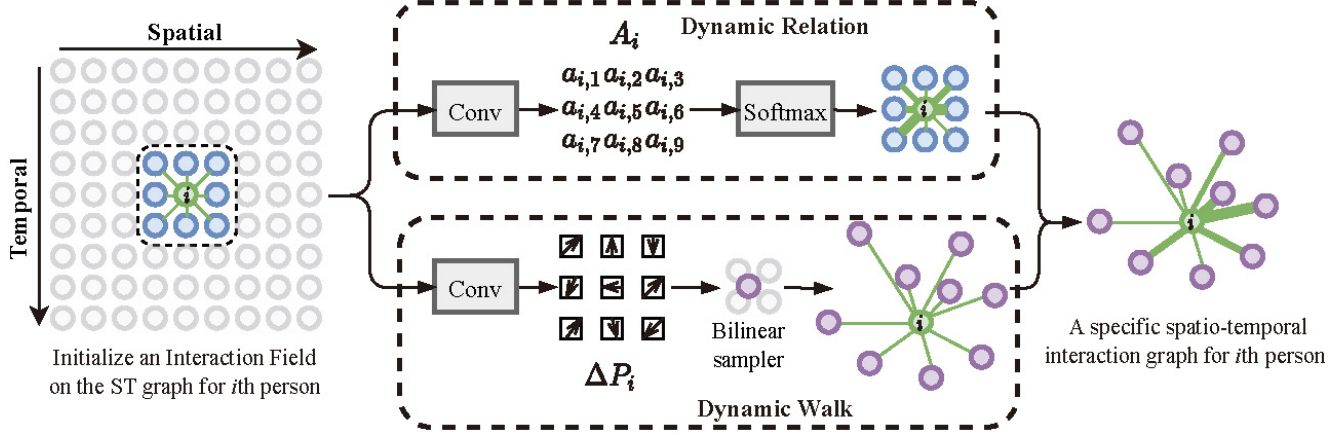


Figure 4. **Details of DR and DW** on creating the person-specific interaction graph for i th person. For the given person, DR predicts a relation matrix and DW predicts the dynamic walk offsets to endow the interaction graph with a global interaction field, both based on an initialized interaction field (we set it to 3×3 as an example).

We perform one-layer ARG to update the person feature as

$$\mathbf{x}_i^{(l+1)} = \sum_{g=1}^{N_g} \sigma \left(\sum_{j=1}^{TN} \tilde{r}_{i,j} \mathbf{x}_j^{(l)} \mathbf{w}^{(g)} \right) + \mathbf{x}_i^{(l)} \quad (3)$$

where N_g, g, l denote the number of graphs in one layer, the graph index and the layer index respectively; σ is an activation function (ReLU in our implementation); $\mathbf{w}^{(g)} \in \mathbb{R}^{D \times D}$ is the graph-specific trainable transformation matrix. Note that $\mathbf{w}^{(g)}$ and $\tilde{r}_{i,j}$ are also layer-specific but for the purpose of clarity, we omit this superscript l . Similar thing is done for learnable parameters and relations in following equations.

After the feature updating, we finally perform a global pooling operation on the reshaped $\mathbf{X}^{(l+1)} \in \mathbb{R}^{T \times N \times D}$ to get the final group representation $\mathbf{z} \in \mathbb{R}^D$. The Cross Inference Block proposed in [46] ameliorates the fully-connected inference by criss-cross inference as shown in Figure 2(b).

3.3. Dynamic Relation

Before we dive into the proposed modules, we present the definition for the interaction field. The **Interaction Field** is a region upon an ST graph that is involved in inferring the interaction features. One example of the interaction field is shown using a dashed box on the ST graph in Figure 4. Our proposed modules, *i.e.* DR and DW, jointly process features within this field to infer person-specific interaction graphs. The initialized interaction field covers a selected person’s spatio-temporal neighborhoods, which provides direct interaction cues. More complex initializations are left for future exploration.

We propose Dynamic Relation (DR) to infer the relation matrix for the person-specific interaction graph. An illustration of DR is shown in the upper branch of Figure 4. ‘Dynamic’ in DR refers to the fact that the relation matrix is

dependent on the features in the initialized interaction field, rather than sticks to the same when updating every feature.

To infer the dynamic relations within this field, we adopt convolution following [11, 50]. For a selected i th feature on the original ST graph, we denote $\mathbf{u}_i \in \mathbb{R}^{(K \times D)}$ as the stacking features within its interaction field and denote K as the interaction field size, *e.g.*, $K = 9$ if the interaction field is 3×3 . We rewrite the convolution in a matrix form as

$$\mathbf{A}_i = \mathbf{W}_a \mathbf{u}_i + \mathbf{b}_a \quad (4)$$

where $\mathbf{W}_a \in \mathbb{R}^{K \times (K \times D)}$ is the linear projection matrix for inferring relations; $\mathbf{b}_a \in \mathbb{R}^K$ is the bias parameters. $\mathbf{A}_i = \{a_{i,k} | k = 1, \dots, K\}$ with $a_{i,k} \in \mathbb{R}^1$ is the relation matrix for i th feature, where k enumerates K features in i th feature’s initialized field. Similar to Eq.2, $\tilde{a}_{i,k}$ is the normalized $a_{i,k}$ along the index k , *i.e.*, $\tilde{a}_{i,k} = \text{softmax}_k(a_{i,k})$.

Instead of updating the features in a fully-connected graph or criss-cross graph, we update the features within the initialized field as

$$\mathbf{x}_i^{(l+1)} = \sigma \left(\sum_{k=1}^K \tilde{a}_{i,k} \mathbf{x}_k^{(l)} \mathbf{w} \right) + \mathbf{x}_i^{(l)} \quad (5)$$

Note that we do not incorporate multiple graphs due to its excessive parameters and trivial improvement [43].

3.4. Dynamic Walk

Although DR has successfully inferred their relations with all person features in the initialized interaction field, it still follows a predefined message passing route, which lacks the ability for person-specific interaction modeling. Moreover, previous methods manage to model long-range spatio-temporal dependency by a fixed graph in a fully-connected or criss-cross scheme, which consumes excessive computational resources. We propose a Dynamic Walk (DW) module that enables features within the interaction

field to execute dynamic walks on the primary ST graph. An illustration of DW is shown in the lower branch of Figure 4. Through DW, we hope to model complex spatio-temporal dependency using a size-limited interaction field. ‘Dynamic’ in DW refers to the fact that the interaction graph is dependent on the features in the initialized interaction field, which is not predefined anymore.

To allow for dynamic walk, we need to predict their spatio-temporal dynamic walk offsets. For a selected i th person feature, we denote the dynamic walk offsets for all features within the interaction field as $\Delta P_i = \{\Delta p_{i,k} | k = 1, \dots, K\}$ with $\Delta p_{i,k} \in \mathbb{R}^2$. We predict the dynamic walk offsets as

$$\Delta P_i = W_p u_i + b_p \quad (6)$$

where $W_p \in \mathbb{R}^{(K \times 2) \times (K \times D)}$ is the linear projection matrix for predicting dynamic walk offsets; $b_p \in \mathbb{R}^{(K \times 2)}$ is the bias term. Similar to DR, this predicts the dynamic walk offsets for all features within the field and it is instantiated by convolution. Using the predicted offsets, we can obtain the dynamic-walked features by performing dynamic walk on the ST graph. Note that the dynamic-walked features are clamped to be within the range of the ST graph. As the dynamic walk offsets are constantly fractional, a bilinear sampler [23] is adopted to sample dynamic-walked features. We denote the coordinate of the k th feature in the i th interaction field as $p_{i,k} \in \mathbb{R}^2$. Dynamic-walked features $Y_i = \{y_{i,k} | k = 1, \dots, K\}$ with $y_{i,k} \in \mathbb{R}^D$ can be formulated as

$$y_{i,k} = \sum_{m=1}^T \sum_{n=1}^N x_{(m-1)N+n} \delta(m, n, p_{i,k}, \Delta p_{i,k}) \quad (7)$$

$$\delta(m, n, p_{i,k}, \Delta p_{i,k}) = \max(0, 1 - |m - p_{i,k}^T - \Delta p_{i,k}^T|) \times \max(0, 1 - |n - p_{i,k}^N - \Delta p_{i,k}^N|) \quad (8)$$

where superscript ‘ T ’ and ‘ N ’ denotes the temporal and spatial coordinate.

Based on the dynamic-walked features, we can update the i th feature as

$$x_i^{(l+1)} = \sigma \left(\sum_{k=1}^K \tilde{a}_{i,k} y_{i,k}^{(l)} w \right) + x_i^{(l)} \quad (9)$$

Note that in the above formulation, we combine the DR and DW to form the final dynamic updating function.

4. Experiments

In this section, we first present datasets and implementation details for GAR. Next, we perform quantitative analysis to explore the contributions of our modules and variances of different interaction field initializations, and to prove the superiority in terms of computational complexity. Then, we compare our methods with previous state-of-the-art methods. Finally, we provide visualizations to understand DIN better.

4.1. Experiment Settings

Datasets So far, there are two widely used datasets in group activity recognition, namely Volleyball dataset (VD) [22] and Collective Activity dataset (CAD) [9].

The Volleyball dataset comprises of 3,493 training clips and 1,337 testing clips, which are trimmed from 55 videos of volleyball matches. For each short clip, it provides three kinds of annotations: i) coordinates of players’ bounding boxes in the center frame of a given clip; ii) individual action labels for the annotated person: *blocking, digging, falling, jumping, moving, setting, spiking, standing* and *waiting*, which are not used in our experiments; iii) group activity labels for the given clip: *right set, right spike, right pass, right winpoint, left set, left spike, left pass* and *left winpoint*. To perform feature extraction on the whole clip, we use the tracklets provided by [4]. Two metrics are used for evaluating the performance of a model, *i.e.*, MCA (%) which is short for Multi-class Classification Accuracy and MPCA (%) which is short for Mean Per Class Accuracy.

The Collective Activity dataset comprises of 44 videos containing varying number of frames from 194 to 1,814 frames. Similar to VD, it is labelled with three levels of annotations: i) coordinates of people’s bounding boxes on the center frame of every ten frames; ii) individual action labels for the annotated person: *NA, crossing, waiting, queueing, walking* and *talking*, which are not used in our experiments; iii) group activity labels for every ten frames: *crossing, waiting, queueing, walking* and *talking*. We follow [42, 45, 46] to merge the class *crossing* and *walking* into *moving*. Similarly, we use the tracklets from [4]. Train-test split follows [32]. MPCA is used for evaluation on this dataset due to class imbalance.

Implementation Details For VD, we use video images with resolution $H \times W = 720 \times 1280$. For CAD, we use video images with resolution $H \times W = 480 \times 720$. For both datasets, we use video clips which contain $T = 10$ frames each following [43, 46, 32, 21]. The maximum number of people in the scene is $N = 12$ for VD and $N = 13$ for CAD. We use person feature with embedding dimension $D = 1024$. For Lite DIN, we use the embedding dimension $D_l = 128$. The convolution operations for DR and DW are initialized by zero vectors [11]. When applying convolution on the graph, we use zero paddings to maintain a fixed interaction field size. We follow [43] to initialize the backbone of DIN model with parameters from the base model. We do not use any action label supervision. For the training of VD, we employ Adam optimizer whose learning rate starts with 1×10^{-4} and decay rate is $\frac{1}{3}$ every 10 epochs. For the training of CAD, we employ the same optimizer whose learning rate starts and stays with 5×10^{-5} . We run 30 epochs in total. The hyper-parameter for Adam is $\beta_1 = 0.9, \beta_2 = 0.999$ and $\epsilon = 10^{-8}$.

Model	MCA	MPCA
Base model	87.8	88.4
DIN w/ DR	92.1	92.3
DIN w/ DW	92.0	92.5
DIN w/ DR+DW	93.1	93.3
DIN w/ DR+DW*	92.9	93.1

Table 1. Ablation study on the usage of **DR and DW**. Experiments are conducted on VD. The backbone is set to ResNet-18.

4.2. Quantitative Analysis

In this subsection, we conduct experiments on VD. We set the backbone for quantitative analysis to ResNet-18.

DR and DW We first conduct ablation study to demonstrate the efficacy of proposed modules. We use a fixed initialized interaction field of 3×3 and following models:

- **Base model**: It consists of a backbone network, RoIAlign, the global pooling layer and a final classification layer.
- **DIN w/ DR**: It contains a backbone, RoIAlign, DR module, a global pooling layer and a classification layer. It allows for a relation matrix prediction within the interaction field.
- **DIN w/ DW**: It contains a backbone, RoIAlign, DW module, a global pooling layer and a classification layer. It allows for a dynamic walk prediction to expand its interaction field.
- **DIN w/ DR+DW**: It is defined analogously with above models. It allows for a dynamic relation prediction based on the original features in the field as illustrated in Figure 4.
- **DIN w/ DR+DW***: It is defined analogously with DIN w/ DR+DW, except that it allows for a dynamic relation prediction based on dynamic-walked features, *i.e.*, Y_i .

The results for the above models are shown in Table 1. The table indicates that incorporating any proposed modules can significantly improve the performance. Compared to ARG [43], the result for DIN w/ DR indicates the superiority of joint processing and a small interaction field. The result for DIN w/ DW indicates the superiority of person-specific interaction graphs. We find that DIN w/ DR and DIN w/ DW show similar improvements compared to the base model. We consider it is because that after we perform the dynamic walk on a graph, the bilinear sampler defines the feature interpolation by bilinear weights, which is, to some extent, one kind of dynamic relations. These weights are determined by the dynamic walk offsets, which are not as straightforward as DR. The partial ability of dynamic relations and a global interaction graph enable DIN w/ DW to perform similarly with DIN w/ DR.

Combining DR and DW, the DIN model is endowed with more dynamicity and larger interaction fields, thus perform-

Reasoning Module	#Params	FLOPs	MCA	MPCA
PCTDM [45]	26.235M	6.298G	90.3	90.5
ARG [43]	25.182M	5.436G	91.1	91.4
AT [15]	5.245M	1.260G	90.0	90.2
HiGCIN[46]	1.051M	184.992G	91.4	92.0
SACRF [31]	29.422M	76.757G	90.7	91.0
EDP	3.146M	0.755G	91.6	91.6
DR	1.140M	0.272G	92.1	92.3
DW	1.222M	0.291G	92.0	92.5
DR+DW	1.305M	0.311G	93.1	93.3
Lite DR+DW	0.180M	0.042G	92.6	92.8

Table 2. **Computational complexity analysis**. Their backbones are set to ResNet-18. #Params and FLOPs for the backbone and the embedding layer are **not** included.

ing even better. Specifically, DIN w/ DR+DW performs slightly better than DIN w/ DR+DW*, which indicates the initialized interaction field provides sufficient information for predicting relations.

Computational Complexity Analysis In this subsection, we present the parameters and FLOPs that the reasoning module contains. Note that the reasoning module that we define does not include the backbone and the person feature embedding layer, as we mainly focus on an efficient reasoning module in this paper. Since previous methods’ reported results vary in the input modality, backbones and implementation details, we re-implement them to fit into our framework and codebase while ensuring consistency to the original paper and their publicly available codes. We use an initialized interaction field of 3×3 for all our proposed modules. For a fair comparison, we set all their backbones to ResNet-18. The results are listed in Table 2. Besides, we present the statistics for the backbone and person feature embedding as a reference: 24.8M #Params, 674.6 GFLOPs for 720×1280 resolution and 24.8M #Params, 254.9 GFLOPs for 480×720 resolution. The result has shown that the model with DR or DW alone achieves higher performances than previous methods, at the same time, reduces the computational complexity of the reasoning module. By combining DR and DW, our model benefits from the dynamicity of person-specific relation matrices and the dynamic walk offsets, thus achieving better performance. Note that the model with DR+DW adds very little computation cost compared to the model with DR or DW alone. The lite model with DR+DW achieves impressive results while adding very little computational overhead compared to the base model.

Besides the proposed model variants, we further provide another reasoning module:

- **EDP**: The corresponding model is DIN w/ EDP. It is similar to DIN w/ DR, except that it uses the Embedded Dot-Product (EDP, formulated in Eq.1) within the interaction field rather than DR for inferring the rela-

Module	Field	#Params	Complexity	FLOPs	Complexity	MCA	MPCA
DR+DW	3×3	1.305M		0.311G		93.1	93.3
	5×5	2.976M	$\Theta(D(3K^2 + D))$	0.712G	$\Theta(TND)$	92.7	93.1
	7×7	8.432M		2.021G	$\Theta(3K^2 + D)$	92.4	92.7
	9×9	21.212M		5.089G		92.5	93.0
ST factorised DR+DW	1×3, 3×1	2.160M		0.516G		92.6	93.0
	1×5, 5×1	2.258M	$\Theta(2D(3K + D))$	0.540G	$\Theta(2TND)$	92.3	92.8
	1×7, 7×1	2.406M		0.575G	$\Theta(3K + D)$	92.8	93.0
	1×9, 9×1	2.602M		0.622G		92.1	92.5
Lite DR+DW	3×3	0.180M		0.042G		92.6	92.8
	5×5	0.387M		0.092G	$\Theta(TND_l)$	92.6	93.1
	7×7	1.069M	$\Theta(D_l(3K^2 + D_l + D))$	0.256G	$\Theta(3K^2 + D_l)$	92.3	92.7
	9×9	2.667M		0.639G	$\Theta(D_l + D)$	92.3	92.5

Table 3. Results for **increasing initialized interaction fields** using three models. Backbone: ResNet-18. Computational cost for the backbone and embedding layer is **not** included. $K = 9, 25, 49, 81$ for 4 interaction fields and $D_l = \frac{D}{8}$.

tion matrix A_j . We set $D_\tau = D$.

We can observe that **i)** the previous pairwise interaction model EDP which predicts relations using only two persons, performs slightly worse than DR and cost higher computation overhead; **ii)** If comparing EDP to ARG, it shows a small initialized field ameliorates the over-smoothing that ARG has due to fully-connected inference, and achieves better performances.

Initialized Interaction Field for DIN and its Variants

To model the spatio-temporal interactions among people, an interaction field with appropriate size should be selected. We mainly provide with experiments on DIN and its two variants to choose an appropriate size: **i)** single interaction field that is initialized to cover a certain spatio-temporal domain, *e.g.*, 3×3 ; **ii)** stacking layers that separately cover spatial and temporal domain (ST factorised model), *e.g.*, 1×3 and 3×1 ; **iii)** lite model that covers a certain spatio-temporal domain. The results for increasing interaction field sizes are shown in Table 3. It indicates that **i)** Larger interaction field sizes will not result in a good performance. **ii)** Similarly, stacking layers to separately model spatial and temporal interactions also result in slightly worse results. We consider they are due to the over-smoothing problem [27] caused by stacking layers or too dense connections, which brings about excessive similarity between person features. **iii)** ST factorised and lite model both distinctly reduce the cost by reducing the exponent of K and the value of D to D_l , while both maintaining better results than previous methods.

4.3. Comparisons with the State-of-the-Art

In this subsection, we compare our models with previous state-of-the-art models. For a fair comparison, we only adopt RGB images as our model input and adopt a single backbone.

Performance on Volleyball dataset The result is shown in Table 4. Generally, our methods can achieve impressive results on this dataset. For methods using ResNet-18, our method can surpass them by 1.7%. For methods us-

Method	Backbone	MCA	MPCA
SBGAR [28]	Inception-v3	66.9	67.6
SSU [4]	Inception-v3	89.9	-
CERN-2 [35]	VGG-16	83.3	83.6
SPA+KD [40]	VGG-16	89.3	89.0
PCTDM [45]	ResNet-18	90.3	90.5
stagNet [32]	VGG-16	89.3	-
CRM [3]	I3D	92.1	-
ARG [43]	ResNet-18	91.1	91.4
PRL [21]	VGG-16	91.4	91.8
AT [15]	ResNet-18	90.0	90.2
SACRF [31]	ResNet-18	90.7	91.0
STBiP* [48]	Inception-v3	91.3	-
HiGCIN [46]	ResNet-18	91.4	92.0
Ours-DIN	VGG-16	93.6	93.8
	ResNet-18	93.1	93.3
Ours-Lite DIN	VGG-16	93.2	93.4
	ResNet-18	92.6	92.8

Table 4. Comparisons with previous state-of-the-art methods on **Volleyball dataset**. We mark with ‘-’ if results are not provided. * denotes results without visual context for fair comparison.

Method	Backbone	MPCA
HDTM[22]	AlexNet	89.7
CERN-2[35]	VGG-16	88.3
Recurrent Modeling[42]	VGG-16	89.4
PCTDM[45]	AlexNet	92.2
stagNet[32]	VGG-16	89.1
SPA+KD[40]	VGG-16	92.5
ARG[43]	ResNet-18	92.3
PRL[21]	VGG-16	93.8
HiGCIN[46]	ResNet-18	93.0
Ours-DIN	VGG-16	95.9
	ResNet-18	95.3
Ours-Lite DIN	VGG-16	94.0
	ResNet-18	93.8

Table 5. Comparisons with previous state-of-the-art methods on **Collective Activity dataset**.

ing VGG-16, our methods can surpass them by 2.2%. If considering the computational overhead of DR and DW, our models show more superiority. Our models generally cost less computational overhead and perform better than RNN-based models like [45, 32, 4, 28, 35], due to a better representation from jointly modeling of spatial-temporal interaction. Our models outperform GNN-based methods like [43, 46, 15], which ascribes to the dynamicity of the proposed modules. Method [43] even uses 16 graphs for reasoning but still trails our model.

Performance on Collective Activity dataset The result is shown in Table 5. With our proposed modules, the model with a VGG-16 backbone outperforms other methods by 2.1% and a ResNet-18 backbone by 2.3%. Note that Lite DIN which costs little extra computation can already

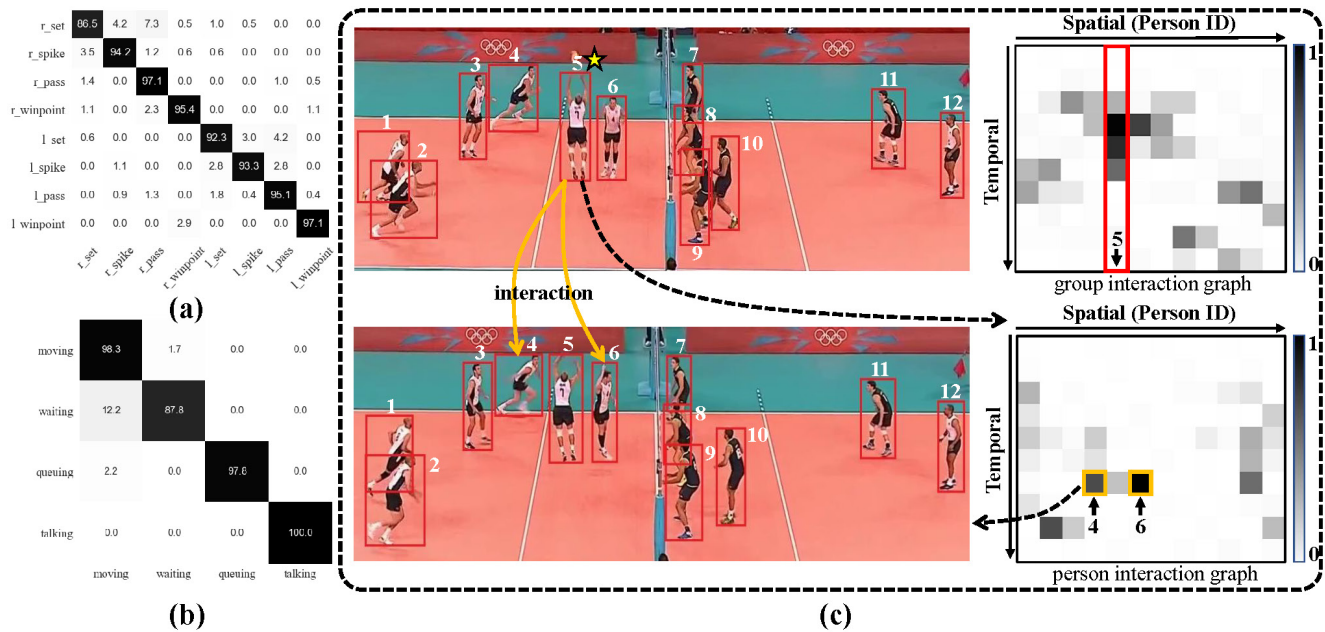


Figure 5. (a) The confusion matrix for Volleyball dataset using VGG-16 as a backbone. (b) The confusion matrix for Collective Activity dataset using VGG-16 as a backbone. (c) Visualizations of a left set activity example. The upper left image is the starting image of the video clip. The upper right is the corresponding group interaction graph. The lower right is the interaction graph of the 5th person (key person, the red box in the group interaction graph). The lower left illustrates two of the 5th person's key interactions (yellow boxes in the 5th person's interaction graph).

achieve results on par with previous best methods, thus revealing the merits of introducing dynamicity.

Confusion Matrices The confusion matrices of VGG-16 models on VD and CAD are respectively shown in Figure 5(a) and (b). For VD, the modeling of dynamic spatial long-range interactions enables the model to distinguish left activities from right activities. Compared with confusion matrices from methods [32, 46], our methods performs well for *pass* and *set* activities. We ascribe it to the dynamic interaction modeling between spatio-temporal persons, because *pass* and *set* activities involve a person passing the ball and a person catching the ball. For CAD, compared with confusion matrices of methods [22, 42, 46], our methods distinguish the *waiting* well. Previous methods mistake *waiting* for *moving* a lot because they fail to distinguish the temporal variations of people, which we tackle well.

4.4. Qualitative Analysis

Group interaction graph First, we visualize the group interaction graph for one example in the upper right image of Figure 5(c), which sums all person-specific interaction graphs. It shows the people whom others interact more with to form the activity. If we sum along the temporal axis, we can find a key person (5th person, the red box in the group interaction graph) with the highest weight. In this example, it is the person performing *setting* action, which is significant in the *left set* group activity.

Person interaction graph We take a step further by visualizing the key person's interaction graph in the lower right image of Figure 5(c), which sums his interaction graphs in different temporal steps (T graphs in total). It indicates that our modules enable global-level interactions though we initialize the interaction field locally. As shown in the person interaction graph, the yellow boxes are two of key interactions with the key person. In this example, they might spike the ball set from the key person.

5. Conclusion and Future Works

In this paper, we propose the Dynamic Inference Network to address the problems of inference on a predefined graph and inference in a computationally expensive way. With limited computation overhead, our model can achieve competitive results on publicly available datasets. Experiments have shown that person-specific interaction context is effective in inferring group activities. More challenging tasks and efficient inference models are left for future exploration. Moreover, this paper focus on the reasoning of person features, while a decent dynamic model incorporating visual context [48] are left for future exploration.

Acknowledgement: We would like to thank Jiayang Ren, Rong Jin and anonymous reviewers for their valuable feedback. This work was supported by the National Science Foundation China grant No. U1609213.

References

- [1] Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [2] Mohamed Rabie Amer, Peng Lei, and Sinisa Todorovic. Hirc: Hierarchical random field for collective activity recognition in videos. In *European Conference on Computer Vision*, pages 572–585. Springer, 2014.
- [3] Sina Mokhtarzadeh Azar, Mina Ghadimi Atigh, Ahmad Nickabadi, and Alexandre Alahi. Convolutional relational machine for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7892–7901, 2019.
- [4] Timur Bagautdinov, Alexandre Alahi, François Fleuret, Pascal Fua, and Silvio Savarese. Social scene understanding: End-to-end multi-person action localization and collective activity recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4315–4324, 2017.
- [5] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 4509–4517, 2016.
- [6] Michael B Chang, Tomer Ullman, Antonio Torralba, and Joshua B Tenenbaum. A compositional object-based approach to learning physical dynamics. *arXiv preprint arXiv:1612.00341*, 2016.
- [7] Tianshui Chen, Weihao Yu, Riquan Chen, and Liang Lin. Knowledge-embedded routing network for scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6163–6171, 2019.
- [8] Wongun Choi and Silvio Savarese. A unified framework for multi-target tracking and collective activity recognition. In *European Conference on Computer Vision*, pages 215–230. Springer, 2012.
- [9] Wongun Choi, Khuram Shahid, and Silvio Savarese. What are they doing?: Collective activity classification using spatio-temporal relationship among people. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1282–1289. IEEE, 2009.
- [10] Wongun Choi, Khuram Shahid, and Silvio Savarese. Learning context for collective activity recognition. In *CVPR 2011*, pages 3273–3280. IEEE, 2011.
- [11] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017.
- [12] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 1, pages 886–893. Ieee, 2005.
- [13] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019.
- [14] Chen Gao, Jiarui Xu, Yuliang Zou, and Jia-Bin Huang. Drg: Dual relation graph for human-object interaction detection. In *European Conference on Computer Vision*, pages 696–712. Springer, 2020.
- [15] Kirill Gavrilyuk, Ryan Sanford, Mehrsan Javan, and Cees GM Snoek. Actor-transformers for group activity recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 839–848, 2020.
- [16] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018.
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Yedid Hoshen. Vain: Attentional multi-agent predictive modeling. *arXiv preprint arXiv:1706.06122*, 2017.
- [20] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [21] Guyue Hu, Bo Cui, Yuan He, and Shan Yu. Progressive relation learning for group activity recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 980–989, 2020.
- [22] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1971–1980, 2016.
- [23] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.
- [24] Tian Lan, Leonid Sigal, and Greg Mori. Social roles in hierarchical models for human activity recognition. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1354–1361. IEEE, 2012.
- [25] Tian Lan, Yang Wang, Weilong Yang, Stephen N Robi-novitch, and Greg Mori. Discriminative latent models for recognizing contextual group activities. *IEEE transactions on pattern analysis and machine intelligence*, 34(8):1549–1562, 2011.
- [26] Jiachen Li, Fan Yang, Masayoshi Tomizuka, and Chiho Choi. Evolvegraph: Multi-agent trajectory prediction with dynamic relational reasoning. *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 2020.

- [27] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [28] Xin Li and Mooi Choo Chuah. Sbgar: Semantics based group activity recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 2876–2885, 2017.
- [29] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, and Cewu Lu. Transferable interactiveness knowledge for human-object interaction detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3585–3594, 2019.
- [30] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, and Charles Leiserson. Evolvegc: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5363–5370, 2020.
- [31] Rizard Renanda Adhi Pramono, Yie Tarng Chen, and Wen Hsien Fang. Empowering relational network by self-attention augmented conditional random fields for group activity recognition. In *European Conference on Computer Vision*, pages 71–90. Springer, 2020.
- [32] Mengshi Qi, Jie Qin, Annan Li, Yunhong Wang, Jiebo Luo, and Luc Van Gool. stagnet: An attentive semantic rnn for group activity recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 101–117, 2018.
- [33] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–417, 2018.
- [34] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofighi, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1349–1358, 2019.
- [35] Tianmin Shu, Sinisa Todorovic, and Song-Chun Zhu. Cern: confidence-energy recurrent network for group activity recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5523–5531, 2017.
- [36] Xiangbo Shu, Jinhui Tang, Guojun Qi, Wei Liu, and Jian Yang. Hierarchical long short-term concurrent memory for human interaction recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [37] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [38] Chen Sun, Per Karlsson, Jiajun Wu, Joshua B Tenenbaum, and Kevin Murphy. Stochastic prediction of multi-agent interactions from partial observations. *arXiv preprint arXiv:1902.09641*, 2019.
- [39] Jinhui Tang, Xiangbo Shu, Rui Yan, and Liyan Zhang. Coherence constrained graph lstm for group activity recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [40] Yansong Tang, Zian Wang, Peiyang Li, Jiwen Lu, Ming Yang, and Jie Zhou. Mining semantics-preserving attention for group activity recognition. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1283–1291, 2018.
- [41] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4601–4607. IEEE, 2018.
- [42] Minsi Wang, Bingbing Ni, and Xiaokang Yang. Recurrent modeling of interaction context for collective activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3048–3056, 2017.
- [43] Jianchao Wu, Limin Wang, Li Wang, Jie Guo, and Gangshan Wu. Learning actor relation graphs for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9964–9974, 2019.
- [44] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5410–5419, 2017.
- [45] Rui Yan, Jinhui Tang, Xiangbo Shu, Zechao Li, and Qi Tian. Participation-contributed temporal dynamic model for group activity recognition. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1292–1300, 2018.
- [46] Rui Yan, Lingxi Xie, Jinhui Tang, Xiangbo Shu, and Qi Tian. Hgcin: Hierarchical graph-based cross inference network for group activity recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [47] Rui Yan, Lingxi Xie, Jinhui Tang, Xiangbo Shu, and Qi Tian. Social adaptive module for weakly-supervised group activity recognition. In *European Conference on Computer Vision*, pages 208–224. Springer, 2020.
- [48] Hangjie Yuan and Dong Ni. Learning visual context for group activity recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3261–3269, 2021.
- [49] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018.
- [50] Li Zhang, Dan Xu, Anurag Arnab, and Philip HS Torr. Dynamic graph message passing networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3726–3735, 2020.
- [51] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9308–9316, 2019.