# Group Activity Prediction with Sequential Relational Anticipation Model

Junwen Chen<sup>[0000-0001-9808-1520]</sup>, Wentao Bao<sup>[0000-0003-2571-3341]</sup>, and Yu Kong<sup>[0000-0001-6271-4082]</sup>

Golisano College of Computing and Information Sciences Rochester Institute of Technology {jc1088,wb6219,yu.kong}@rit.edu

**Abstract.** In this paper, we propose a novel approach to predict group activities given the beginning frames with incomplete activity executions. Existing action<sup>1</sup> prediction approaches learn to enhance the representation power of the partial observation<sup>2</sup>. However, for group activity prediction, the relation evolution of people's activity and their positions over time is an important cue for predicting group activity. To this end, we propose a sequential relational anticipation model (SRAM) that summarizes the relational dynamics in the partial observation and progressively anticipates the group representations with rich discriminative information. Our model explicitly anticipates both activity features and positions by two graph auto-encoders, aiming to learn a discriminative group representation for group activity prediction. Experimental results on two popularly used datasets demonstrate that our approach significantly outperforms the state-of-the-art activity prediction methods.

**Keywords:** Activity prediction, Group activity, Structured prediction, Relational model.

## 1 Introduction

Group activity prediction is to forecast an activity performed by a group of people before the activity ends. Different from group activity recognition, it only has access to the beginning frames of a video containing **incomplete** activity execution. It is useful in the scenarios where the intelligent systems have to make prompt decisions, such as surveillance and traffic accident avoidance where multiple people are present. Unfortunately, existing action prediction methods [52,51,23,22,21] are limited to actions performed by an individual. Even though some methods [23,24,45] attempt to predict actions performed by multiple people in standard databases such as UCF101 [39], they simply model the multiple people as a single entity and ignore their relations. This would undoubtedly result in a low prediction performance.

We define action as the behavior performed by a single person, and define activity as the behavior performed by a group of people.
 We define *partial observation* as the beginning frames with incomplete activity execution, and

<sup>&</sup>lt;sup>2</sup> We define *partial observation* as the beginning frames with incomplete activity execution, and *full observation* as the one with complete activity execution.



Fig. 1: Given the beginning frames, our method models the relational dynamics of a group, and predicts a group activity by anticipating the group activity representation and their positions occurring in the future unobserved frames.

As shown in [23], one of the major challenges in activity prediction is how to enhance the discriminative power of the features extracted from the partial observations. However, this is even more challenging to do so in group activity prediction as multiple people are present in the scene. Each person's individual action may vary and people's interactions frequently appear and change in a group activity. To this end, it is important to model the relations of multiple people in the observed frames and predict their future group representations. In addition, if only limited beginning frames are observed, it would be extremely difficult to directly anticipate the features of full observations at once. A temporally progressive anticipation model is desired for modeling activity evolution.

To address these challenges, we propose a novel sequential relational anticipation model (SRAM) for group activity prediction by anticipating group activities and positions in the future (see Fig. 1). SRAM is developed as an encoder-decoder framework, in which an observation encoder summarizes the relational dynamics in the beginning observed frames and a sequential decoder further anticipates the representations for group activities and positions occurring in the future. Specifically, the observation encoder naturally models the relational dynamics of people and complex interactions between people in the observed frames. To predict group activity, we propose a sequential decoder to anticipate the structured group representation in the future using several unrolling stages. Two graph auto-encoders are used in the sequential decoder to explicitly anticipate the activity and the position relations of people in the unobserved frames. We propose to make sequential prediction that *progressively* anticipates the future group representation by performing multiple unrolling stages guided by three novel loss functions. This allows us to better capture complex group activity evolution.

To our best knowledge, we are the first to investigate the challenging problem of group activity prediction. The benefit of our method is twofold. Firstly, it not only predicts people's group activities but also predicts individuals' positions in the future. Our experimental results show that predicting people's future positions significantly helps predict their group activities. Secondly, the proposed method progressively anticipates structured group representations, which has shown to be very powerful in prediction especially when limited frames are observed. This idea could be generalized to other prediction tasks, e.g. human motion prediction [31] and video prediction [47].

Our contributions can be summarized as follows:

- We propose a novel sequential decoder to anticipate the representations for multiple people's future positions and activity, aiming to learn a discriminative structural representation for group activity prediction.
- We progressively anticipate the structured group representations at several unrolling stages guided by novel loss functions. This improves the performance when only few frames are observed.
- Extensive experiments demonstrate that our method outperforms the existing state-of-the-arts by a large margin.

## 2 Related work

Action Prediction aims to recognize the label of an action before the action is fully executed. Existing work [26,6,21,17,36,30,53,37,43] focuses on predicting actions performed by an individual. Ryoo et al. [34] used integral and dynamic bag-of-words to represent features variations over time. DeepSCN [23] and AAP-Net [24] make use of sequential context information by transferring knowledge in full videos to partial observations. Wang et al. [45] developed a teacher-student learning framework to distill knowledge from the action recognition task, in order to enhance action prediction. Gammulle et al. [11] presented a jointly learnt task for both action prediction and future motion representation inference.

Prediction on interactions between two people was studied in [51,52]. Yan et al. [51] developed a tri-coupled recurrent structure and an attention mechanism to address action prediction for two individuals' interactions. Yao et al. [52] predicted the motion of the interactions between two people, but did not predict their interaction labels. Different from them, we focus on the prediction of group activities involving multiple people. Our method elegantly captures complex relational dynamics between people for learning discriminative information.

**Group Activity Recognition** has been extensively studied in previous work [2,38,19,49]. Early work applies graphical models on the extracted handcraft features [2,27,46] as group representations. Deep learning methods for multi-people activity recognition have shown excellent performance [38,4,44,18,8,41,12]. HDTM [19] develops a two-stage LSTM model to firstly extract features of temporal individual motions and then aggregate neighborhood information. SSU [4] achieves the individual detection and group activity recognition in a unified framework. Recent work suggests that only part of people's motions contribute to the entire group activity [49,10,33,3,16], via suppressing the irrelevant actions. Previous work also shows that interactions between people are important in understanding group activity. For example, HRN[18] introduces a hierarchical spatial relational layer to learn the relational representations between two players. Other methods, including Stagnet [32], S-RNN [5], SBGAR [29] apply



Fig. 2: Overall architecture. Our framework SRAM takes the beginning  $t_0$  observed frames as input and predicts the group activity label. An observation encoder first summarizes the relational dynamics in partial observation as a latent variable  $Z_0$ . Then, a sequential decoder takes over  $Z_0$  and progressively anticipates the group representation through K unrolling stages. The output of the last unrolling stage is expected to contain rich discriminative information for group activity prediction. Details can be seen in Fig. 3.

structural-RNN to obtain spatiotemporal features. ARG [48] explicitly models the interactions by employing graph convolution on a learnable graph.

The main difference between our work and group activity recognition methods is that we aim at predicting the group activity label given *incomplete* activity execution, while these methods are given complete activity executions. This prompts us to develop novel model architecture and loss functions in this work.

## 3 Our approach

4

J. Chen et al.

**Problem formulation.** Our goal is to predict the activity label y of a group of people given a partial observation of a video containing incomplete activity execution. We define the *observation ratio* as the number of observed frames  $t_0$  in a streaming video divided by the total number of frames T in the corresponding full video following [23,24], i.e.  $t_0/T$ . For instance, if a partial video contains 30 frames and the corresponding full video contains 100 frames, then the observation ratio of this activity is 30%.

During training, we have access to all full training videos containing complete group activity executions. These full videos are supposed to contain all the discriminative information for classification. During test, given a partial observation of a group activity execution, we encourage our model to anticipate the group representations that contain similar amount of discriminative information as the corresponding full observation. Thus, its prediction power can be enhanced.

**Overall architecture.** The overall architecture is shown in Fig. 2. We formulate our group activity prediction model as an encoder-decoder framework that contains an *observation encoder* and a *sequential decoder*. Given a partial observation containing  $t_0$  frames, the observation encoder summarizes the relational dynamics of the group from the partial observation and then the sequential decoder anticipates the group representation for activities and positions in the future unobserved frames.

Due to the large motion variations between a partial and a full observation, a novel sequential decoder is proposed in this work to progressively anticipate the structured group representation for the future unobserved frames by several unrolling stages. This is useful for enhancing the discriminative power of the anticipated representation, especially if limited frames are observed. Moreover, for group activity, relations between multiple people are discriminative information and they vary as time. To predict group activity, our sequential decoder uses two graph auto-encoders to concurrently perform relational anticipation on both people's activity features and their positions.

### 3.1 Relation Modeling for Group Activity

Given  $t_0$  observed frames, we first extract features of all the observed  $t_0$  frames, and then apply ROIAlign [15] to extract the feature vectors of multiple people based on their positions  $\{B_1, B_2, \dots, B_{t_0}\}(t \in \{1, \dots, t_0\})$ . Action features and positions of the *i*-th individual on the *t*-th frame are represented as  $\mathbf{x}_t(i)$  and  $\mathbf{b}_t(i)$  respectively. Afterwards, upon the individual dynamics, we follow [48] to explicitly model the pair-wise position relations and action relations of multiple people in the observed frames as two relation graphs  $G_t^a \in \mathbb{R}^{N \times N}$  and  $G_t^p \in \mathbb{R}^{N \times N}$ , respectively. Both of the two graphs have N nodes representing N people in the *t*-th frame. Given the *i*-th and *j*-th individuals, the edge of the action similarity graph  $G_t^a(i, j)$  is computed by the cosine similarity and normalized by Softmax function. The edge on the position relation graph  $G_t^p(i, j)$  is computed by the normalized Euclidean distance (denoted by  $d(\cdot, \cdot)$ ):

$$G_{t}^{\mathbf{a}}(i,j) = \frac{\exp\left(\mathbf{x}_{t}(i)^{\mathrm{T}} \cdot \mathbf{x}_{t}(j)\right)}{\sum_{j=1}^{N} \exp\left(\mathbf{x}_{t}(i)^{\mathrm{T}} \cdot \mathbf{x}_{t}(j)\right)}, \ G_{t}^{\mathrm{p}}(i,j) = \frac{1/d(\mathbf{b}_{t}(i),\mathbf{b}_{t}(j))}{\sum_{j=1}^{N} 1/d(\mathbf{b}_{t}(i),\mathbf{b}_{t}(j))}.$$
 (1)

Once the graphs are built, we obtain the structured representations for the group activity in the observed frames. We will also perform anticipation on the two graphs representing the group activity in the unobserved frames, which will be discussed below.

#### 3.2 Observation Encoder $\mathcal{E}$

The observation encoder  $\mathcal{E}$  is proposed to summarize spatiotemporal information of the complex relational dynamics of multiple people in partial observations containing  $t_0$  frames.  $\mathcal{E}$  learns to map  $G_{1:t_0}^{a}$ ,  $G_{1:t_0}^{p}$ , and  $X_{1:t_0}$  to a latent variable  $Z_0$ , by the spatio-temporal graph convolution network ST-GCN [50]. Specifically, it first performs spatial graph convolution [20] on the two graphs  $G_t^{p}$  and  $G_t^{a}$  for the t-th frame

$$\sigma(G_t^{\rm p} X_t W_{\rm p}) + \sigma(G_t^{\rm a} X_t W_{\rm a}), \tag{2}$$

and then performing temporal convolution [28] on every three consecutive frames to learn the latent variable  $Z_0$ . Here,  $\sigma$  is ReLU activation,  $W_p$  and  $W_a$  are learnable weights, and  $X_t$  is the action features of N people. Latent variable  $Z_0$ will be integrated in the sequential decoder, and guides its unrolling stages.

6 J. Chen et al.



Fig. 3: Sequential decoder  $\mathcal{D}$  is formulated as two auto-encoders  $\mathcal{E}_{a}$ - $\mathcal{D}_{a}$  and  $\mathcal{E}_{p}$ - $\mathcal{D}_{p}$  that progressively anticipate the group activity representation for future unobserved frames using multiple unrolling stages. At the *k*-th stage,  $\mathcal{D}$  is fed with the summary of the partial observation encoded in the latent variable  $Z_{0}$  as well as the action features  $\hat{X}_{k}$  and the position features  $\hat{B}_{k}$  from the previous stage. Then  $\mathcal{D}$  anticipates the action features  $\hat{X}_{k+1}$  and positions  $\hat{B}_{k+1}$ .

Different from ARG [48], our model captures the temporal patterns of people's relations, which is useful for group activity prediction.

### 3.3 Sequential Decoder $\mathcal{D}$

The performance of state-of-the-art action prediction methods [24,23] is still limited especially when few beginning frames are given. This is mainly because they use a direct mapping from partial observation to the corresponding full observation in one pass, which is not powerful enough to deal with large visual variations between partial and full observations. In this paper, we propose a sequential decoder that *progressively* anticipates the group representation that is expected to contain rich discriminative information as the fully observed activity using K unrolling stages (see Fig. 3). This allows us to create a more powerful model for group activity prediction.

Besides, different from individual action prediction methods [24,45], people's relations formulated as graphs using Eq. (1), are discriminative information for group activity. Moreover, the group activity varies overtime. It is necessary to predict group representations by anticipating relations in the unobserved stage. As described in Section 3.1, people's relations can be inferred from their action similarity and relative positions. For example, a partial observation of a volleyball activity is given, which contains run-up of ace spikers and waiting gestures of their opponents. Our model is supposed to predict it as "spiking" by the cue that the players are moving towards net with their actions. Therefore, we develop a sequential decoder as a mixture of two graph auto-encoders: an activity auto-encoder  $\mathcal{E}_{\rm a}$ - $\mathcal{D}_{\rm a}$  for predicting activity representations and a position auto-encoder are coupled by the shared latent variables  $Z_0$  learned from partial observations.

Activity auto-encoder  $\mathcal{E}_{\mathbf{a}}$ - $\mathcal{D}_{\mathbf{a}}$ . Using K activity auto-encoders, the proposed sequential decoder progressively anticipates the activity representation by K unrolling stages. Each activity auto-encoder at the k-th stage ( $k \in \{1, 2, \dots, K\}$ ) is fed with the output  $\hat{X}_k$  of the activity auto-encoder at the previous (k-1)-th stage. We use the spatiotemporal action features at the last observed frame  $t_0$  as the input of the activity auto-encoder at stage k = 1. We encode the input  $\hat{X}_k$  of current unrolling stage to a latent variable  $Z_k^a$  by

$$Z_k^a = \sigma(G_k^p \hat{X}_k U_{ep}) + \sigma(G_k^a \hat{X}_k U_{ea}), \qquad (3)$$

and then decodes the activity representation  $\hat{X}_{k+1}$ :

$$\hat{X}_{k+1} = \sigma(G_k^{\rm a}(Z_0 + Z_k^{\rm a}))U_{\rm da}) + \sigma(G_k^{\rm p}(Z_0 + Z_k^{\rm a})U_{\rm dp}),\tag{4}$$

where  $U_{ep}$ ,  $U_{ea}$ ,  $U_{dp}$ ,  $U_{da}$  are learnable parameters.  $\hat{X}_{k+1}$  is the anticipated group features at the k-th stage, and is served as the input for the activity auto-encoder at the (k+1)-th stage. The anticipation of  $\hat{X}_{k+1}$  is conditioned on latent variables  $Z_k^a$  and  $Z_0$ , in order to both keep track of the short-term information of the previous unrolling stage and use the long-term spatiotemporal information in the partial observations.  $G_k^a$  and  $G_k^p$  are computed by the generated activity features and positions at the k-th stage using similar functions as Eq. (1), respectively (replacing time step t by the stage k).

The benefits of the progressive anticipation using K unrolling stages lie in two aspects. First, the temporal dependency of activity evolution is naturally built between successive stages. This allows us to naturally anticipate structured group activity representations for prediction purpose. Second, the prediction granularity can be controlled with the number of unrolling stages K. The case when K = 1 is equivalent to the existing one-pass solution used in [24,23].

**Position auto-encoder**  $\mathcal{E}_{\mathbf{p}}$ - $\mathcal{D}_{\mathbf{p}}$ . As described in Section 3.1, the interactions between two people also depend on their relative positions. Thus, it is necessary to explicitly anticipate the positions of these people in group activity prediction.

Similar to activity auto-encoder, the proposed sequential decoder also performs K unrolling stages for position prediction for a group of people using Kposition auto-encoders. Each position auto-encoder at stage k is fed with the output of its previous auto-encoder at stage k - 1, and outputs the positions of people. Experimental results in Section 4.4 show that the anticipated future positions of people help improve performance of group activity prediction.

The position auto-encoder first encodes the positions  $B_k$  of multiple people to a latent variable  $Z_k^p$  at stage k through graph convolution [20]:

$$Z_k^{\rm p} = \sigma(G_k^{\rm p}\hat{B}_k V_{\rm ep}) + \sigma(G_k^{\rm a}\hat{B}_k V_{\rm ea}), \tag{5}$$

and then decodes the positions  $\hat{B}_{k+1}$  for the next stage by

$$\hat{B}_{k+1} = \sigma(G_k^{\rm a}(Z_0 + Z_k^{\rm p}))V_{\rm da}) + \sigma(G_k^{\rm p}(Z_0 + Z_k^{\rm p})V_{\rm dp}),\tag{6}$$

where  $V_{ep}$ ,  $V_{ea}$ ,  $V_{dp}$ ,  $V_{da}$  are learnable parameters.  $G_k^p$  and  $G_k^a$  are the same graphs used in the activity auto-encoder. The anticipation of  $\hat{B}_{k+1}$  is conditioned on

latent variables  $Z_k^{\rm p}$  and  $Z_0$ , in order to both keep track of the short-term position information of the previous unrolling stage and use the long-term spatiotemporal information in the partial observations.

Position prediction is also benefited by sequential prediction via several unrolling stages, since the prediction granularity can be controlled. Similar to the activity auto-encoder, the position auto-encoder at stage k = 1 also takes the positions  $B_{t_0}$  of people on the last observed frame as input. The activity auto-encoder and the position auto-encoder share the same graphs  $G_k^{\rm p}$  and  $G_k^{\rm a}$  and are both conditioned on the latent variable  $Z_0$  (see Fig. 3).

#### 3.4 Feature Aggregation for Prediction

SRAM returns both group activity and position representations at each of K unrolling stages. The K-th stage corresponds to the full observation status, which contains the most discriminative information of an activity. We disregard all the outputs given by the activity autoencoders from the 1-st to (K - 1)-th stages, and perform max-pooling on the output  $\hat{X}_{K+1}$  given by the activity autoencoder at the K-th stage as the group activity representations. The resulting feature vector is used for group activity prediction. Similarly, we directly use the output  $\hat{B}_{K+1}$  given by the K-th position autoencoder to perform position prediction.

### 3.5 Loss Functions and Model Learning

Adversarial loss. Inspired by [13], we encourage SRAM to generate representations corresponding to ground-truth full observations. We use two discriminators for the features generated by the sequential decoder. Discriminator  $D_1$ is an activity classifier implemented by a softmax layer.  $\mathcal{L}_{cls}$  is computed on the output of  $D_1$ . Discriminator  $D_2$  is an adversarial regularizer and tells the difference between the generated group features  $\hat{X}_{1:K}$  and group features of full videos  $F_{1:K}(X)$ . Using the adversarial loss, SRAM is encouraged to generate features that are indistinguishable from the group features of the corresponding full videos:

$$\mathcal{L}_{\text{GAN}} = \mathbb{E}_{X_{(1:T)} \sim p_{\text{data}}(X_{(1:T)})} \log D_2 \left( F_{1:K}(X) \right)$$

$$+ \mathbb{E}_{X_{(1:t_0)} \sim p_{\text{data}}(X_{(1:t_0)})} \log \left( 1 - D_2(\hat{X}_{1:K}) \right).$$
(7)

Note that the generated group representation  $\hat{X}_{1:K}$  is computed by SRAM S from the partial observation  $X_{1:t_0}$ .

**Sequential reconstruction loss** is proposed to align the predicted activity representations to become close to the ground-truth activity representations at each unrolling stage. Since our method has K-stage sequential prediction, it is necessary to encourage the predicted group representations  $\hat{X}_{1:K}$  on each of the K unrolling stages to become close to the ground-truth features at that timestamp. This is different from adversarial loss that only align the generated features to be close to ground-truth at full observation stage.

We train a separate ST-GCN  $F(\cdot)$  as a recognition model to obtain the group activity representations of full videos X for training. The resulting frame-wise group representations F(X) are used to encourage the activity features of the *i*th person generated at the *k*-th unrolling stage to be similar to the ground-truth features using

$$\mathcal{L}_{\rm rec} = \frac{1}{K \times N} \sum_{k=1}^{K} \sum_{i=1}^{N} \left\| \hat{\mathbf{x}}_{k}(i) - F_{k}(X, i) \right\|^{2}.$$
 (8)

Here,  $F_k(X, i)$  is the features of the *i*-th person of the full video at the *k*-th stage. This loss function sequentially computes the difference between the predicted features  $\hat{\mathbf{x}}_k(i)$  (the *i*-th row on  $\hat{X}_{1:K}$ ) for the *i*-th person at unrolling stage *k* and the ground-truth features  $F_k(X, i)$ , mimicking how a partial observation is progressively approaching its corresponding full observation.

**Position regression loss.** We use the tracklets of individuals provided by [18] as the ground-truth of individuals positions. During training, we use the mean square error between the predicted positions and ground-truth positions at the K unrolling stages as loss function:

$$\mathcal{L}_{\text{reg}} = \frac{1}{K \times N} \sum_{k=1}^{K} \sum_{i=1}^{N} ||\hat{\mathbf{b}}_{k}(i) - \mathbf{b}_{k}(i)||^{2},$$
(9)

where the predicted position  $\hat{\mathbf{b}}_k(i)$  is the *i*-th row of  $\hat{B}_k$ , i.e., the *i*-th person's position predicted by the sequential decoder at the *k*-th stage.

Model learning. During training, the overall objective function is written as a sum of sequential reconstruction loss  $\mathcal{L}_{rec}$ , adversarial loss  $\mathcal{L}_{GAN}$ , classification loss  $\mathcal{L}_{cls}$  implemented by softmax loss, and position regression loss  $\mathcal{L}_{reg}$ :

$$\min_{\mathcal{E},\mathcal{D}} \max_{D_1,D_2} \mathcal{L}_{\rm rec} + \mathcal{L}_{\rm GAN} + \mathcal{L}_{\rm cls} + \mathcal{L}_{\rm reg}.$$
 (10)

Sequential relational anticipation model  $(\mathcal{E}, \mathcal{D})$  and two discriminators  $(D_1, D_2)$  are alternatively trained until convergence.

#### 3.6 Discussion

**Group activity modeling and anticipation.** Our SRAM captures the interactions of multiple people in the observation encoder, and anticipates their future relations by a sequential decoder. This is different from existing action prediction methods [24,32] that can only predict the action of an individual. We believe such a novel method will pave the way for future research in other structured visual prediction.

**Structured sequential prediction.** Compared with group activity recognition methods [48,32,18], our method performs *sequential prediction* of group activity, in form of future positions and activity representations. Our activity prediction is also facilitated by explicitly predicting people's future positions.

Activity evolution over time. Our sequential decoder progressively predicts future representations through several unrolling stages, which boosts performance when only few frames are observed. It is guided by a sequential reconstruction loss, mimicking how a partial observation is sequentially approaching its full observation and an adversarial loss to make the generated full observation features to become indistinguishable from the real full observation features.

#### Experiments 4

#### 4.1 Datasets

**Volleyball Dataset** [19] consists of 4830 video clips distributed in 8 group activities, such as *left spiking* and *right setting*. Each clip has 41 frames. [19] provides the players' tracklets and splits the dataset into training, validation and testing sets. Existing group activity recognition methods [48,19,18,32] use the middle 10 frames of each video. To generalize it to prediction task, we extend it to use the middle 20 frames as full observations, in order to model sequential dynamics. Note that the middle 20 frames contain complete group activity executions, because athletes generally move quickly to complete a group activity. such as direct spiking in a volleyball game.

Collective Activity Dataset (CAD) [7] contains 44 videos with 5 group activities, including crossing, queueing, walking, talking and waiting. The group activities in CAD are labeled as the majority of people's individual actions. We use the existing tracklet information and training/testing splits following [48]. The number of the frames in videos ranges from 100 to 2000. Following [32,48,4], we divide each video into 10-frame video clips. This expands training and testing data to 1746 and 765 clips, respectively. CAD mainly contains periodic activities such as *walking*, in which significant changes can be seen in 10 frames.

#### 4.2**Implementation Details**

Following [48], we extract a 1024-dimensional feature vector for each individual with tracklets provided by [19], using Inception-v3 [40] as backbone and ROIAlign [15]. We use three steps for training: First, Inception-v3 pretrained on ImageNet is fine-tuned on single frames by jointly predicting individual actions and group activities. Then, we freeze the backbone and finetune the recognition model  $F(\cdot)$  given full videos in the training set. The recognition model contain two ST-GCN layers [50], both with 256-d hidden units. After that, we train the proposed model. The observation encoder has two layers ST-GCN with both 256-d hidden units. The activity auto-encoder's encoder has one graph convolution layer that encodes the input into 256-d latent feature space. The position auto-encoder has two layer graph convolution by encoding the 2-d positions into 64-d space and then 256-d latent space. During training, SRAM plus classifier  $D_1$  and discriminator  $D_2$  are alternatively updated.

The experiments are conducted with 10 different observation ratios ranging from 10% to 100% of full videos length. The number of unrolling stages K is set

10

11

to 5. We use stochastic gradient descent for optimization. For Volleyball dataset, the three steps are trained for 30 epochs, 10 epochs and 20 epochs with learning rate 0.001, 0.001, 0.0001 respectively. For Collective Activity Dataset, the three steps are trained for 20 epochs, 50 epochs and 10 epochs with learning rate 0.0001, 0.0001, 0.0005, respectively.

### 4.3 Comparison with State-of-the-art

We compare our method with the state-of-the-art prediction methods LRCN [42], DeepSCN [23], IBoW and DBoW [34], KD [45], AAPNet [24] and state-of-the-art group activity recognition methods, including HRN [18], HDTM [19] ARG [48], SSU [4]. Following these methods' original setting, LRCN and HDTM adopt the AlexNet [25] as the backbone. HRN, IBow, DBow, DeepSCN and original AAPNet use VGG-19. Our method follows ARG and SSU to use Inception-V3 method. HRN, HDTM, ARG, SSU and our method adopt the tracklets of players provided by [19]. To make a fair comparison, we extend state-of-the-art action prediction method AAPNet ("e-AAPNet" for simplification) to make use of tracking information and use Inception-V3 as backbone. We train all of the comparison methods using the parameters described in their original papers.

**Results on Volleyball dataset.** Table 1 summarizes the prediction performance of the proposed method, existing action prediction methods and group activity recognition methods. Results demonstrate that our model outperforms the comparison methods. Existing action prediction methods, e.g., LRCN, IBoW, DBoW, DeepSCN, AAPNet, KD propose to improve the prediction performance by information transfer. However, they regard multiple people as a single entity and do not consider the interactions between multiple people. Thus, the extracted features do not contain informative cues of the interactions of people, resulting in a low prediction performance. The proposed method uses tracklets [19], while the existing predictors for individuals e.g. LRCN, IBoW, DBoW, DeepSCN, AAPNet, KD do not. To make a fair comparison, we extend AAPNet to use tracklet information. Experimental results show that our method can predict the dynamics of interactions and better enrich partial observations.

Group activity recognition methods such as HDTM, SSU, HRN, ARG do not have capability of gaining extra information from full activity executions. Thus, when the observation ratio is very low (10% or 20% observations), their performance is much lower than our method. Note that ARG applies random sampling strategy by sampling three frames from an entire video as input. In the comparison experiment, this strategy is applied in each of the partial observations as input. The proposed method consistently outperforms ARG, as our method captures the temporal dynamics of multiple people in the group, and sequentially generates features close to the corresponding full observations. It improves the representation power of the partial observations, and facilitates group activity prediction.

Table 1: Group activity prediction accuracy (%) on Volleyball dataset with observation ratios ranging from 10% to 100%. Group activity recognition results can be seen from the last column, in which 100% frames are observed.

					0/					0/		
Models	Tracklet	Backbones	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
LRCN [9]	No	AlexNet	48.17	51.61	54.67	57.44	59.76	61.23	63.75	64.32	64.77	65.37
HDTM [19]	Yes	AlexNet	52.43	59.09	66.04	76.37	80.48	81.82	84.07	84.47	84.60	84.06
IBoW [35]	No	VGG	58.03	60.72	64.84	65.26	67.51	70.80	73.45	74.24	74.29	75.63
DBoW [35]	No	VGG	58.03	55.56	56.16	58.93	59.90	61.97	63.79	63.06	63.88	64.78
DeepSCN [23]	No	VGG	59.46	62.23	65.52	70.38	72.55	77.37	79.75	80.35	80.31	80.78
HRN [18]	Yes	VGG	52.58	56.99	64.32	74.49	76.96	80.36	83.72	84.74	84.08	85.30
KD [45]	No	VGG	65.67	67.68	70.00	70.83	71.96	72.10	73.22	73.30	73.30	73.90
AAPNet [24]	No	VGG	59.53	65.37	68.29	72.25	75.24	77.79	79.91	80.25	80.18	80.78
e-AAPNet [24]	Yes	InceptionV3	62.98	70.31	77.64	83.55	84.91	85.86	87.54	87.23	87.92	89.01
SSU [4]	Yes	InceptionV3	63.20	70.65	79.66	84.07	87.13	87.65	88.30	88.18	88.41	89.01
ARG [48]	Yes	InceptionV3	64.82	69.41	76.07	79.43	82.70	83.99	85.04	85.19	85.86	85.94
Ours	Yes	InceptionV3	77.86	82.57	84.97	87.06	88.63	88.93	89.08	88.93	88.48	91.97

**Results on Collective Activity Dataset.** 

Comparison results are listed in Table 2. Our method outperforms existing methods ARG, DeepSCN, and AAPNet by a large margin. Given tracklets as input, our method is 6.54% higher than e-AAPNet at 50% observation ratio since the people's actions and relations are predicted in our model. Group activities such as group walking are cyclic, and thus the prediction performance of our

Table 2: Prediction accuracy (%) on Collective Activity Dataset.

Models	Tracklet	50%	100%
ARG [48]	Yes	88.10	88.37
DeepSCN [23]	No	81.31	82.22
AAPNet [24]	No	81.57	82.75
e-AAPNet [24]	Yes	86.01	86.67
Ours	Yes	92.55	92.81

method at 50% observation ratio is close to the one at 100% observation ratio.

### 4.4 Ablation Study

We perform detailed ablation studies on the Volleyball dataset to evaluate the contributions of the sequential prediction strategy, as well as the loss functions.

How much does the prediction loss help? The impacts of loss functions are analyzed on Volleyball dataset in detail. The evaluation results can also validate the contributions of the proposed sequential prediction strategy. We compare the following the variants, including: (1) without the position regression loss  $\mathcal{L}_{reg}$  defined in Eq. (9). In this variant, the position auto-encoder for predicting future positions is not used. During sequential prediction, we replace the individuals' positions in the future frames by the ones given by the last observed frame's. The positions are used for computing  $G^{\rm p}$  for each unrolling stage. (2) without adversarial loss  $\mathcal{L}_{\rm GAN}$ . (3) without sequential reconstruction loss  $\mathcal{L}_{\rm rec}$ for generated features of unrolling stages. (4) The proposed full network.

Compared with variant (1), the significant performance gains with all different observation ratios show that the prediction of people's positions is of high importance for group activity prediction. Compared with variants (2) and (3), it shows that the adversarial loss  $\mathcal{L}_{\text{GAN}}$  and the reconstruction loss  $\mathcal{L}_{\text{rec}}$  in our Table 3: Ablation studies on volleyball dataset. We show the accuracy(%) given videos of observation ratio at 10%, 40%, 70%.

(a) Comparison of different loss functions,  $\mathcal{L}_{\text{GAN}}$ , (b) Comparison on the number  $\mathcal{L}_{\text{rec}}$ ,  $\mathcal{L}_{\text{reg}}$ , and  $\mathcal{L}_{\text{cls}}$ .

Loss	10%	40%	70%	Average	K	10%	40%	70%	Average
$(1)\mathcal{L}_{ ext{GAN}}{+}\mathcal{L}_{ ext{rec}}{+}\mathcal{L}_{ ext{cls}}$	75.09	85.59	87.06	84.85	1	70.38	80.02	86.14	81.47
$(2) \; \mathcal{L}_{ m rec} {+} \mathcal{L}_{ m reg} {+} \mathcal{L}_{ m cls}$	76.14	85.79	88.18	86.30	<b>2</b>	72.36	86.59	89.07	85.22
$(3) \mathcal{L}_{ m reg} + \mathcal{L}_{ m GAN} + \mathcal{L}_{ m cls}$	77.61	83.22	85.64	86.22	5	77.86	87.06	89.08	86.85
(4) Ours	77.86	87.67	89.08	86.85	10	77.93	86.69	89.23	86.79

method improve the performance by 0.55% and 0.63% on average, respectively. Therefore, the proposed sequential decoder guided by  $\mathcal{L}_{\text{GAN}}$  and  $\mathcal{L}_{\text{rec}}$  can generate more discriminative activity representations at each stage.

How much does the sequential prediction help? Our sequential decoder predicts group activity representations at K unrolling stages. In this experiment, we evaluate the effect of the number of unrolling stages K on the prediction performance. We set K to 1, 2, 5, and 10, and compare the prediction performance. Table 3b indicates that the best overall prediction performance is achieved when K = 5. The prediction performance is slightly affected when K = 10, but the computational complexity of the prediction model is increased due to the extra unrolling stages. The average prediction performance drops to 81.47% if K = 1. The variant with K = 1 is the one that directly maps partial observation in one unrolling stage, similar to what [24,23] do. The result demonstrates the superiority of our progressive prediction in anticipating discriminative group representations given partial observations. If more stages are allowed (K = 5 or K = 10), the sequential decoder in our model can progressively generate discriminative features for group activity prediction even though it is given very limited frames. Therefore, its prediction performance is improved.

#### 4.5 Position Prediction Evaluation

Visualization of predicted positions As shown in Fig. 4, we visualize the movement of individuals learned by the position auto-encoder in SRAM. The position auto-encoder progressively predicts the positions of individuals at the unrolling stages. The visualization result shows our position auto-encoder can successfully predict the directions and step sizes of individuals in the future based on partial observations. Although Fig. 4 (bottom-right) shows the direction of the predicted movement is mostly accurate, the future position of a person is not accurate if the person moves very fast.

**Quantitative evaluation** We quantitatively evaluate our position prediction results compared to two popular trajectory prediction methods SocialGAN [14]



Fig. 4: Visualization of position predictions. The blue and the yellow lines denote the prediction positions and ground-truth positions [19], respectively. " $\times$ " indicates the starting point of the movement. Best viewed in color.

and SocialLSTM [1]. Following SocialGAN, Final Displacement Error (FDE) is used to compute the Euclidean distance between the predicted positions and ground-truth positions at the final timestamp and Average Displacement Error (ADE) is used to compute that at each unrolling stage.

As shown in Tab. 4, the results demonstrate that our method can accurately predict the future positions for a group of people, and our method outperforms the two trajectory prediction methods. This is mainly because we capture the relational action dynamics of multiple people while SocialGAN and SocialLSTM do not.

Table -	4: Fi	nal D	ispla	$\operatorname{cement}$	Error
(FDE)	and	Avera	age	Displac	ement
Error (.	ADE)	for p	ositio	on predi	iction.

Method	FDE	ADE
SocialGAN [14]	5.32	3.05
SocialLSTM [1]	6.44	4.44
Ours	3.62	2.44

## 5 Conclusion

We have proposed a novel sequential relational anticipation model (SRAM) to predict group activity given only the beginning frames of an activity execution. Our model captures complex relational dynamics of multiple people in the observed frames. It then anticipates the group representations including group activity features and position features. A novel sequential decoder is proposed to progressively anticipates the group representations through several unrolling stages. Extensive results on two datasets demonstrate that our method significantly outperforms the state-of-the-art methods. Results also validate that the progressive anticipation using multiple unrolling stages facilitates group activity prediction. Further experimental results show that the modeling and prediction of people's positions improves our performance on group activity prediction.

Acknowledgement: We thank Nvidia for the GPU donation. This research is supported in part by ONR Award N00014-18-1-2875.

## References

- Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social lstm: Human trajectory prediction in crowded spaces. In: CVPR. pp. 961– 971 (2016)
- Amer, M.R., Lei, P., Todorovic, S.: Hirf: Hierarchical random field for collective activity recognition in videos. In: ECCV. pp. 572–585 (2014)
- Azar, S.M., Atigh, M.G., Nickabadi, A., Alahi, A.: Convolutional relational machine for group activity recognition. In: CVPR. pp. 7892–7901 (2019)
- Bagautdinov, T., Alahi, A., Fleuret, F., Fua, P., Savarese, S.: Social scene understanding: End-to-end multi-person action localization and collective activity recognition. In: CVPR. pp. 4315–4324 (2017)
- Biswas, S., Gall, J.: Structural recurrent neural network (srnn) for group activity analysis. In: WACV. pp. 1625–1632 (2018)
- Cai, Y., Li, H., Hu, J.F., Zheng, W.S.: Action knowledge transfer for action prediction with partial videos pp. 8118–8125 (2019)
- Choi, W., Shahid, K., Savarese, S.: What are they doing?: Collective activity classification using spatio-temporal relationship among people. In: ICCV Workshops. pp. 1282–1289 (2009)
- Deng, Z., Vahdat, A., Hu, H., Mori, G.: Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition. In: CVPR. pp. 4772–4781 (2016)
- Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR. pp. 2625–2634 (2015)
- Gammulle, H., Denman, S., Sridharan, S., Fookes, C.: Multi-level sequence gan for group activity recognition. In: Asian Conference on Computer Vision. pp. 331–346 (2018)
- Gammulle, H., Denman, S., Sridharan, S., Fookes, C.: Predicting the future: A jointly learnt model for action anticipation. In: ICCV. pp. 5562–5571 (2019)
- Gavrilyuk, K., Sanford, R., Javan, M., Snoek, C.G.M.: Actor-transformers for group activity recognition. In: CVPR (June 2020)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS. pp. 2672–2680 (2014)
- Gupta, A., Johnson, J., Fei-Fei, L., Savarese, S., Alahi, A.: Social gan: Socially acceptable trajectories with generative adversarial networks. In: CVPR. pp. 2255– 2264 (2018)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. pp. 2961–2969 (2017)
- Hu, G., Cui, B., He, Y., Yu, S.: Progressive relation learning for group activity recognition. CVPR (2020)
- Hu, J.F., Zheng, W.S., Ma, L., Wang, G., Lai, J.H., Zhang, J.: Early action prediction by soft regression. TPAMI (2018)
- Ibrahim, M.S., Mori, G.: Hierarchical relational networks for group activity recognition and retrieval. In: ECCV. pp. 721–736 (2018)
- Ibrahim, M.S., Muralidharan, S., Deng, Z., Vahdat, A., Mori, G.: A hierarchical deep temporal model for group activity recognition. In: CVPR. pp. 1971–1980 (2016)

- 16 J. Chen et al.
- Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
- Kong, Y., Gao, S., Sun, B., Fu, Y.: Action prediction from videos via memorizing hard-to-predict samples. In: AAAI (2018)
- Kong, Y., Kit, D., Fu, Y.: A discriminative model with multiple temporal scales for action prediction. In: ECCV. pp. 596–611 (2014)
- Kong, Y., Tao, Z., Fu, Y.: Deep sequential context networks for action prediction. In: CVPR. pp. 1473–1481 (2017)
- 24. Kong, Y., Tao, Z., Fu, Y.: Adversarial action prediction networks. TPAMI (2020)
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. pp. 1097–1105 (2012)
- Lan, T., Chen, T.C., Savarese, S.: A hierarchical representation for future action prediction. In: ECCV. pp. 689–704 (2014)
- 27. Lan, T., Sigal, L., Mori, G.: Social roles in hierarchical models for human activity recognition. In: CVPR. pp. 1354–1361 (2012)
- Lea, C., Flynn, M.D., Vidal, R., Reiter, A., Hager, G.D.: Temporal convolutional networks for action segmentation and detection. In: CVPR. pp. 156–165 (2017)
- Li, X., Choo Chuah, M.: Sbgar: Semantics based group activity recognition. In: ICCV. pp. 2876–2885 (2017)
- Ma, S., Sigal, L., Sclaroff, S.: Learning activity progression in lstms for activity detection and early detection. In: CVPR. pp. 1942–1950 (2016)
- Martinez, J., Black, M.J., Romero, J.: On human motion prediction using recurrent neural networks. In: CVPR. pp. 2891–2900 (2017)
- Qi, M., Qin, J., Li, A., Wang, Y., Luo, J., Van Gool, L.: stagnet: An attentive semantic rnn for group activity recognition. In: ECCV. pp. 101–117 (2018)
- Ramanathan, V., Huang, J., Abu-El-Haija, S., Gorban, A., Murphy, K., Fei-Fei, L.: Detecting events and key actors in multi-person videos. In: CVPR. pp. 3043–3053 (2016)
- Ryoo, M.S.: Human activity prediction: Early recognition of ongoing activities from streaming videos. In: ICCV. pp. 1036–1043 (2011)
- Ryoo, M., Aggarwal, J.: Stochastic representation and recognition of high-level group activities. IJCV 93(2), 183–200 (2011)
- Sadegh Aliakbarian, M., Sadat Saleh, F., Salzmann, M., Fernando, B., Petersson, L., Andersson, L.: Encouraging lstms to anticipate actions very early. In: ICCV. pp. 280–289 (2017)
- Shi, Y., Fernando, B., Hartley, R.: Action anticipation with rbf kernelized feature mapping rnn. In: ECCV. pp. 301–317 (2018)
- Shu, T., Todorovic, S., Zhu, S.C.: Cern: confidence-energy recurrent network for group activity recognition. In: CVPR. pp. 5523–5531 (2017)
- Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012)
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR. pp. 2818–2826 (2016)
- 41. Tang, Y., Wang, Z., Li, P., Lu, J., Yang, M., Zhou, J.: Mining semantics-preserving attention for group activity recognition. In: Multimedia. pp. 1283–1291 (2018)
- Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: ICCV. pp. 4489–4497 (2015)
- Vondrick, C., Pirsiavash, H., Torralba, A.: Anticipating visual representations from unlabeled video. In: CVPR. pp. 98–106 (2016)
- 44. Wang, M., Ni, B., Yang, X.: Recurrent modeling of interaction context for collective activity recognition. In: CVPR. pp. 3048–3056 (2017)

Group Activity Prediction with Sequential Relational Anticipation Model

- 45. Wang, X., Hu, J.F., Lai, J.H., Zhang, J., Zheng, W.S.: Progressive teacher-student learning for early action prediction. In: CVPR. pp. 3556–3565 (2019)
- 46. Wang, Z., Shi, Q., Shen, C., Van Den Hengel, A.: Bilinear programming for human activity recognition with unknown mrf graphs. In: CVPR. pp. 1690–1697 (2013)
- 47. Wichers, N., Villegas, R., Erhan, D., Lee, H.: Hierarchical long-term video prediction without supervision pp. 6038–6046 (2018)
- Wu, J., Wang, L., Wang, L., Guo, J., Wu, G.: Learning actor relation graphs for group activity recognition. In: CVPR. pp. 9964–9974 (2019)
- Yan, R., Tang, J., Shu, X., Li, Z., Tian, Q.: Participation-contributed temporal dynamic model for group activity recognition. In: Multimedia. pp. 1292–1300 (2018)
- 50. Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: AAAI (2018)
- Yan, Y., Ni, B., Yang, X.: Predicting human interaction via relative attention model. IJCAI pp. 3245–3251 (2017)
- Yao, T., Wang, M., Ni, B., Wei, H., Yang, X.: Multiple granularity group interaction prediction. In: CVPR. pp. 2246–2254 (2018)
- Zhao, H., Wildes, R.P.: Spatiotemporal feature residual propagation for action prediction. In: ICCV. pp. 7003–7012 (2019)