

HIERARCHICAL GRAPH-RNNS FOR ACTION DETECTION OF MULTIPLE ACTIVITIES

Sovan Biswas, Yaser Souri, Juergen Gall

University of Bonn, Germany

ABSTRACT

In this paper, we propose an approach that spatially localizes the activities in a video frame where each person can perform multiple activities at the same time. Our approach takes the temporal scene context as well as the relations of the actions of detected persons into account. While the temporal context is modeled by a temporal recurrent neural network (RNN), the relations of the actions are modeled by a graph RNN. Both networks are trained together and the proposed approach achieves state of the art results on the AVA dataset.

Index Terms— Spatio-temporal action detection, Graph-RNN

1. INTRODUCTION

With the advent of deep neural networks and the availability of large datasets in the last decade, the performance of recent algorithms for action recognition has improved drastically. Annotating large amount of data, however, is very expensive. In particular for spatio-temporal action recognition and localization where multiple actions occur at the same time, action labels and bounding boxes would be required for each frame in order to learn a model using full supervision. The large-scale AVA 2.1 dataset [1] for recognizing and localizing multiple actions in videos, therefore, provides only temporally sparse annotations, i.e., the persons and the actions the people perform are annotated for only one frame per second. This requires to develop methods that can be trained with such sparse annotations [1, 2, 3, 4].

While these works take temporal information into account, they do not model the interactions of the individual persons. While some actions of different persons are uncorrelated, other actions refer to interactions between persons like ‘talk to’ and ‘listen to’. There are also actions that are often performed by several persons at the same time like ‘sit’, ‘stand’, ‘walk’, or ‘play instrument’, but there are also actions that exclude each other. For instance, if a person ‘drives’ a car, it is very unlikely that the other persons in the car ‘stand’ or ‘play an instrument’.

In this work, we therefore propose an approach that learns the relations of actions that occur at the same time and that can be learned using only sparse annotations. In order to

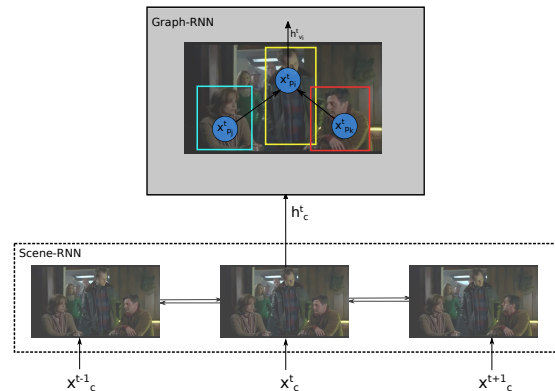


Fig. 1. The proposed network jointly captures temporal context and the relations between different persons by a hierarchy. Scene RNNs, at the lower level, are used to model the temporal context of a scene. The Graph-RNN on top of it, models the relations of the actions of the detected persons.

address the sparseness of the annotations, we do not rely on tracked bounding boxes, which can be unreliable. Instead, we propose a hierarchical model that models the temporal scene context in the lower level and the relations of actions of the detected persons on the top level as illustrated in Fig. 1. The temporal scene context is important since the relations of the actions depend on the scene. For instance, it matters if persons are inside of a moving car or a parking car. We model the scene context by a recurrent neural network (RNN) that uses I3D features [5] as input. The RNN models the temporal context of the entire frames. At the top level of the hierarchy, we combine the hidden states of the scene context RNN with I3D features extracted for all detected persons in a frame. To learn the relations of the actions of all detected persons, we use a graph recurrent neural network [6, 7, 8]. The proposed model therefore learns the scene RNN and the graph RNN together.

We evaluate our approach on the large-scale AVA 2.1 dataset [1] where the approach achieves state-of-the-art results for action detection of multiple activities.

2. RELATED WORK

A common approach for action recognition and localization [9, 10, 11, 12] comprises the detection of the bounding

The work has been funded by the ERC Starting Grant ARCA (677650).




Fig. 2. After detecting bounding boxes using Faster RCNN and pooling I3D features for each detected bounding box, the proposed hierarchical Graph-RNN predicts multiple class labels for each bounding box. In this example, the network detects two persons and infers the activities ‘sit’, ‘watch (a person)’, and ‘listen to (a person)’ for the red bounding box and ‘stand’, ‘talk to (a person)’, and ‘watch (a person)’ for the green bounding box.

boxes in each frame using object detectors [13, 14]. The detected bounding boxes are then linked to obtain action tubes, which are then classified. These approaches, however, assume that every frame is annotated. Since such dense annotations are very time-consuming, the approaches [15, 3, 16] deal with sparse annotations where the action labels and locations are annotated only for a subset of frames, e.g., each frame per second. These works, however, treat each person independently although persons tend to interact with each other.

In the context of group activity analysis [17, 18, 19], the relations between various individual persons are used to infer the action label of the group as well as the individuals. The works [17, 18] propose hierarchical models where the individual actions are modeled at the lower level and the group activity at the top level. Our task, however, differs from group activity analysis. While [17, 18, 19] assume that each individual is part of a sports team or group and each individual performs only one action as part of a group, [1] contains multiple activities per individual and only a subset of the actions are based on interactions with other individuals.

There are several types of networks that can be applied to graphs like Graph Convolutional Networks (GCN) [20, 21] or Graph Recurrent Neural Networks (Graph RNN) [6, 8, 22]. These graph networks have been used in various computer vision applications such as object detection [23], semantic segmentation [24], or visual question answering [25]. For instance, [26] use an attentional GCN to model spatial relations between objects in an image. In our work, we use a Graph RNN combined with a Scene RNN to model relations between different persons as well as temporal context.

3. DETECTION OF MULTIPLE ACTIVITIES

As it is defined in the AVA 2.1 dataset [1], the goal of the approach is to estimate for a frame the bounding boxes of all persons performing an action and for each bounding box the action labels. In contrast to other datasets, each bounding box is annotated by multiple labels. For instance, a person might ‘stand’, ‘carry/hold’ an object, and ‘listen to’ another person

at the same time.

3.1. Features

As illustrated in Fig. 2, the proposed approach consists of two steps. We first detect the persons in a frame and then we use the hierarchical graph RNN to infer the action labels for each detected bounding box. For detecting the bounding boxes, we use Faster RCNN [27] with a ResNet architecture [28]. The detector is fine-tuned on the dataset and the detections are performed in an action class agnostic way.

For the temporal context of a frame t , we extract I3D features [5] for t and its neighboring frames. We only consider the frames t that are annotated in AVA, i.e., one frame per second, but a single I3D feature is computed over a temporal context of 33 frames. This means that the temporal receptive fields of the extracted features for the annotated frames t and $t + 1$ slightly overlap. Given l annotated frames before and after frame t , we get a sequence of $L = 2l + 1$ I3D features

$$x_c^{t-l}, \dots, x_c^{t-1}, x_c^t, x_c^{t+1}, \dots, x_c^{t+l}, \quad (1)$$

which corresponds to a temporal context of approximately L seconds for AVA.

We also use the detected bounding boxes for each sparsely sampled frame and use region pooling to extract I3D features per detected bounding box as in [1]. This gives for each detected bounding box i in frame t another I3D feature vector that is denoted by $x_{p_i}^t$. We will use these features as input for the hierarchical graph RNN.

3.2. Hierarchical Graph RNN (HGRNN)

The proposed hierarchical graph RNN as illustrated in Fig. 1 comprises two types of RNNs that are trained together. At the lower level, the scene RNN, which is described in Section 3.2.1, models the temporal scene context. The scene context estimated by the scene RNN is then used as input for the graph RNN, which models the relations of the actions performed by the detected persons. The graph RNN, which is described in Section 3.2.2, predicts then for each bounding box multiple action labels.

3.2.1. Scene RNN

The scene RNN takes all L scene features x_c^t as well as the features extracted for each detected person $x_{p_i}^t$ in all L frames as input and predicts a hidden state h_c^t only for frame t . Since it uses the frames before and after the frame t , we use a bidirectional RNN with GRUs. At this stage, we do not model any relations between the persons and simply perform max-pooling over all features $x_{p_i}^t$ at each frame.

$$x_P^t = \text{maxpool}_i(x_{p_i}^t) \quad (2)$$

$$x^t = x_c^t \oplus x_P^t$$

$$h_c^t = \text{biGRU}(x^t; h_c^{t-1}, h_c^{t+1}). \quad (3)$$

While x_P^t denotes the maxpooled person feature for all detected persons in frame t , \oplus denotes the concatenation of two vectors.

3.2.2. Graph RNN

Given i detected persons in frame t and the corresponding features $x_{p_i}^t$ as well as the hidden scene representation h_c^t estimated by the scene RNN, we now model the relations of the detected persons to infer the activities for each of them. To this end, we represent each detected person as a node ($v_i \in V$) of a fully connected graph. This means that we consider all possible relations how an action of a person effects the actions of the other persons.

As in [6, 7, 8], we use a graph RNN that iteratively updates the hidden representation for each node v_i based on the intermediate representations of the other nodes. In our case, the equations for the graph RNN are given by

$$\begin{aligned} x_{v_i}^{(j)} &= \text{maxpool}_{v_i \in V}(h_{v_i}^{(j-1)}) \\ a_{v_i}^{(j)} &= x_{p_i} \oplus h_c \oplus x_{v_i}^{(j)} \\ z_{v_i}^{(j)} &= \sigma(U^z a_{v_i}^{(j)} + W^z h_{v_i}^{(j-1)}) \\ r_{v_i}^{(j)} &= \sigma(U^r a_{v_i}^{(j)} + W^r h_{v_i}^{(j-1)}) \\ s &= \tanh(U^s a_{v_i}^{(j)} + W^s (h_{v_i}^{(j-1)} \circ r_{v_i}^{(j)})) \\ h_{v_i}^{(j)} &= (1 - z_{v_i}^{(j)}) \circ h_{v_i}^{(j-1)} + z_{v_i}^{(j)} \circ s \end{aligned} \quad (4)$$

where we omitted the frame index t for the ease of reading and \circ denotes the Hadamard product. At each iteration j , the hidden representation for a detected person is given by $h_{v_i}^{(j)}$. To update the representation, we first maxpool the hidden representation over all nodes and concatenate it with the original person feature x_{p_i} as well as the temporal scene context h_c estimated by the scene RNN, which provides a longer temporal context for the graph RNN than the person features x_{p_i} . Using a GRU variant, $h_{v_i}^{(j)}$ is then updated. After a fix number of iterations, the estimated hidden representation for each detected person $h_{v_i}^{(j)}$ is then fed to a fully connected layer with sigmoid as activation function to infer all action classes that are simultaneously performed.

Temporal context l	0	1	3
mAP	19.0%	19.5%	20.9%

Table 1. Impact of the temporal context l .

The entire hierarchical graph RNN consisting of the scene RNN and the graph RNN is trained jointly using the focal loss [29] for multi-label classification.

3.3. Implementation Details

The person detector is initialized by a ResNet-101 architecture trained on ImageNet. We finetune the person detector using Adam optimizer with a variable learning rate starting from 0.00001 and dropping by 0.5 at regular intervals. This fine-tuning is done for 100K steps with an effective batch size of 20. The I3D network is pre-trained on Kinetics [5]. Due to memory reasons, we reduce the size of the I3D features from 1024 dimensions to 256 dimensions using a fully connected layer on top of the I3D network and finetune the network on the dataset. For finetuning, we use Adam optimizer with a variable learning rate starting from 0.0005 and dropping by 0.5 at various intervals. This fine-tuning is done for 50K steps with an effective batch size of 50. To make our model robust, we perform data augmentation using random flips and crops as in [1].

The hierarchical graph RNN is randomly initialized and trained from scratch using 1000 steps with a batch size of 50 using Adam optimizer with a constant learning rate of 0.0001. Furthermore, we use the focal loss [29] with $\gamma = 2$. We train the network on the annotated ground-truth bounding boxes. For inference, we use the detected bounding boxes and we perform a simple multiplication of the person detection confidence with the corresponding action prediction score to obtain the final class predictions for each detection. Finally, class specific non-maximum suppression is used to remove duplicate detections.

4. EXPERIMENTS

For evaluation, we use the AVA 2.1 dataset [1]. It contains 60 action classes across 235 videos of 15 minutes each for training and 64 videos of the same length for the validation set, which we use for evaluation. Sparse annotations in form of action labels and bounding boxes are provided for a single frame every second. The evaluation is performed using frame-level mean average precision (frame-AP) at IoU threshold 0.5, as described in [1].

Temporal Context: To analyze the effect of increasing the temporal support, we evaluate various values for the temporal context l (1). Since the frames are sparsely sampled, $l = 3$ corresponds to a temporal context of 7 seconds while $l = 0$ corresponds to 1 second. For the experiment, we use only RGB data without optical flow. The results in Tab. 1

Models	Graph-RNN	Scene-RNN	HGRNNs
mAP	19.0%	19.8%	20.9%

Table 2. Comparison of the Scene-RNN and Graph-RNN with the HGRNN.

Iterations	1	2	3
mAP	20.6%	20.9%	20.8%

Table 3. Impact of the number of iterations.

show that increasing the temporal context increases the accuracy. This is due to the fact that certain actions such as ‘open’ and ‘close’ can be better recognized with a larger temporal receptive field.

Joint Temporal and Interaction Modeling: In Tab. 1, the accuracy for $l = 0$ corresponds to the case where only the Graph-RNN but not the Scene-RNN is used. In Tab. 2, we also report the accuracy if we use only the Scene-RNN but not the Graph-RNN. In both cases, the proposed Hierarchical Graph-RNN, which combines both RNNs in a single model, achieves a higher accuracy. This shows that both the temporal context as well as the interactions between the persons contribute to the action detection accuracy.

Number of Iterations: As discussed in Section 3.2.2, the HGRNN block is iterated. The results in Tab. 3 show that not many iterations are required. This can be attributed to the fact that our Graph-RNN already incorporates temporal information through the hierarchy and it requires only two iterations to update the hidden state of each person based on the hidden states of the other persons. In all other experiments, we use 2 iterations.

Ground Truth (GT) Bounding Boxes: In order to understand the effect of the accuracy of the person detector on the action detection accuracy, we used ground truth bounding boxes during inference. The fine-tuned Faster RCNN person detector achieves an mAP of 89.09% for detecting the annotated bounding boxes on AVA. If ground truth detections are used instead, the action detection accuracy increases by 5 – 6% as shown in Tab. 4. We also report the results if we use RGB and optical flow for computing the I3D features. The additional optical flow increases the accuracy by 2.7% and 3.9% for detected and GT bounding boxes, respectively.

Comparison with State of the Art: The proposed approach outperforms the approaches [1, 2] by a large margin. Most interesting is the comparison to [1] since it uses the same features but a vanilla I3D head for action detection. The proposed hierarchical Graph-RNN improves the accuracy by 6.4% on RGB data and 8.0% on RGB+Flow data. This clearly demonstrates the capability of the proposed hierarchical GRNN in comparison to the I3D head [1]. We also compare our approach with the very recent works [3, 4]. While our approach outperforms [4], [3] achieves a higher accuracy for RGB data. The gain of the accuracy is the use of

Method	GT	Detected
RGB	25.2%	20.9%
RGB+Flow	29.1%	23.6%

Table 4. Quantitative comparison of the proposed method with ground truth bounding boxes and detected bounding boxes.

Method	flow	mAP
AVA [1]		14.5%
ACRN [2]		17.4%
Better AVA [3]		21.9%
HGRNN - RGB		20.9%
AVA [1]	✓	15.6%
D3D [4]	✓	23.0%
HGRNN - Flow	✓	23.6%

Table 5. Comparison of the proposed method with other state of the art methods. A ✓ at the flow column indicates if optical flow has been used.

a single Faster RCNN framework that detects the bounding boxes and the action classes together. Using the proposed hierarchical Graph-RNN within a Faster RCNN framework is therefore a future research direction to improve the accuracy further. However, it is unclear how much gain can be achieved if optical flow is used in addition since [3] does not report any results for optical flow.

5. CONCLUSION

In this paper, we proposed hierarchical Graph Recurrent Neural Networks for recognizing and localizing multiple activities that occur at the same time. The model learns the temporal context as well as the interactions of the detected persons to recognize the actions. In our experimental evaluation, we have shown that the proposed model outperforms a temporal as well as a graph RNN and that the proposed approach achieves state of the art results on the AVA dataset.

6. REFERENCES

- [1] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al., “AVA: A Video dataset of Spatio-Temporally Localized Atomic Visual Actions,” in *CVPR*, 2018, pp. 6047–6056.
- [2] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid, “Actor-Centric Relation Network,” in *ECCV*, 2018, pp. 318–334.

- [3] Rohit Girdhar, João Carreira, Carl Doersch, and Andrew Zisserman, “A better baseline for AVA,” *arXiv preprint arXiv:1807.10066*, 2018.
- [4] Jonathan C Stroud, David A Ross, Chen Sun, Jia Deng, and Rahul Sukthankar, “D3d: Distilled 3d networks for video action recognition,” *arXiv preprint arXiv:1812.08249*, 2018.
- [5] Joao Carreira and Andrew Zisserman, “Quo Vadis, Action recognition? A new model and the kinetics dataset,” in *CVPR*, 2017, pp. 4724–4733.
- [6] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini, “The Graph Neural Network Model,” *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 61–80, 2009.
- [7] Marco Gori, Gabriele Monfardini, and Franco Scarselli, “A New Model for Learning in Graph Domains,” in *IJCNN*, 2005, pp. 729–734.
- [8] Yujia Li, Richard Zemel, Marc Brockschmidt, and Daniel Tarlow, “Gated Graph Sequence Neural Networks,” in *ICLR*, 2016.
- [9] Georgia Gkioxari and Jitendra Malik, “Finding action tubes,” in *CVPR*, 2015, pp. 759–768.
- [10] Rui Hou, Chen Chen, and Mubarak Shah, “Tube Convolutional Neural Network (T-CNN) for Action Detection in Videos,” in *ICCV*, 2017, pp. 5822–5831.
- [11] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid, “Action Tubelet Detector for spatio-temporal action localization,” in *ICCV*, 2017, pp. 4415–4423.
- [12] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip Torr, and Fabio Cuzzolin, “Online Real-Time multiple spatiotemporal action localisation and prediction,” in *ICCV*, 2017, pp. 3657–3666.
- [13] Ross Girshick, “Fast R-CNN,” in *ICCV*, 2015, pp. 1440–1448.
- [14] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar, “Rethinking the Faster R-CNN Architecture for Temporal Action Localization,” in *CVPR*, 2018, pp. 1130–1139.
- [15] Philippe Weinzaepfel, Xavier Martin, and Cordelia Schmid, “Towards weakly-supervised action localization,” *arXiv preprint arXiv:1605.05197*, 2016.
- [16] Dong Li, Zhaofan Qiu, Qi Dai, Ting Yao, and Tao Mei, “Recurrent tubelet proposal and recognition networks for action detection,” in *ECCV*, 2018, pp. 303–318.
- [17] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori, “A Hierarchical Deep Temporal model for Group Activity Recognition,” in *CVPR*, 2016, pp. 1971–1980.
- [18] Sovan Biswas and Juergen Gall, “Structural Recurrent Neural Network (SRNN) for Group Activity Analysis,” in *WACV*, 2018, pp. 1625–1632.
- [19] Yansong Tang, Zian Wang, Peiyang Li, Jiwen Lu, Ming Yang, and Jie Zhou, “Mining Semantics-Preserving Attention for Group Activity Recognition,” in *ACM-MM*, 2018, pp. 1283–1291.
- [20] Thomas N Kipf and Max Welling, “Semi-supervised classification with Graph Convolutional Networks,” in *ICLR*, 2017.
- [21] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P Adams, “Convolutional Networks on Graphs for Learning Molecular Fingerprints,” in *NIPS*, 2015, pp. 2224–2232.
- [22] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena, “Structural-RNN: Deep Learning on Spatio-Temporal Graphs,” in *CVPR*, 2016, pp. 5308–5317.
- [23] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu, “Learning Human-Object interactions by Graph Parsing Neural Networks,” in *ECCV*, 2018, pp. 401–417.
- [24] Xiaojuan Qi, Renjie Liao, Jiaya Jia, Sanja Fidler, and Raquel Urtasun, “3D Graph Neural Networks for RGBD Semantic Segmentation,” in *CVPR*, 2017, pp. 5199–5208.
- [25] Damien Teney, Lingqiao Liu, and Anton van den Hengel, “Graph-Structured Representations for Visual Question Answering,” in *CVPR*, 2017, pp. 3233–3241.
- [26] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh, “Graph R-CNN for Scene Graph Generation,” in *ECCV*, 2018, pp. 670–685.
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun, “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks,” in *NIPS*, 2015, pp. 91–99.
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep Residual Learning for Image Recognition,” in *CVPR*, 2016, pp. 770–778.
- [29] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar, “Focal Loss for Dense Object Detection,” in *ICCV*, 2017, pp. 2980–2988.