

Spatio-Temporal Action Graph Networks

Roei Herzig^{1*,†}, Elad Levi^{2*}, Huijuan Xu^{3*}, Hang Gao³, Eli Brosh²,
Xiaolong Wang³, Amir Globerson¹, Trevor Darrell^{2,3}

¹Tel Aviv University, ²Nexar, ³UC Berkeley

Abstract

Events defined by the interaction of objects in a scene are often of critical importance; yet important events may have insufficient labeled examples to train a conventional deep model to generalize to future object appearance. Activity recognition models that represent object interactions explicitly have the potential to learn in a more efficient manner than those that represent scenes with global descriptors. We propose a novel inter-object graph representation for activity recognition based on a disentangled graph embedding with direct observation of edge appearance. In contrast to prior efforts, our approach uses explicit appearance for high order relations derived from object-object interaction, formed over regions that are the union of the spatial extent of the constituent objects. We employ a novel factored embedding of the graph structure, disentangling a representation hierarchy formed over spatial dimensions from that found over temporal variation. We demonstrate the effectiveness of our model on the Charades activity recognition benchmark, as well as a new dataset of driving activities focusing on multi-object interactions with near-collision events. Our model offers significantly improved performance compared to baseline approaches without object-graph representations, or with previous graph-based models.

1. Introduction

Recognition of events in natural scenes poses a challenge for deep learning approaches to activity recognition, since an insufficient number of training examples are typically available to learn to generalize to all required observation conditions and variations in appearance. For example, in driving scenarios critical events are often a function of the spatial relationship of prominent objects, yet available event training data may not exhibit variation across a sufficiently wide range of appearances. E.g., if a conventional deep model has only seen red pickup trucks rear-end blue sedans, and green trucks always drive safely in the training set, it

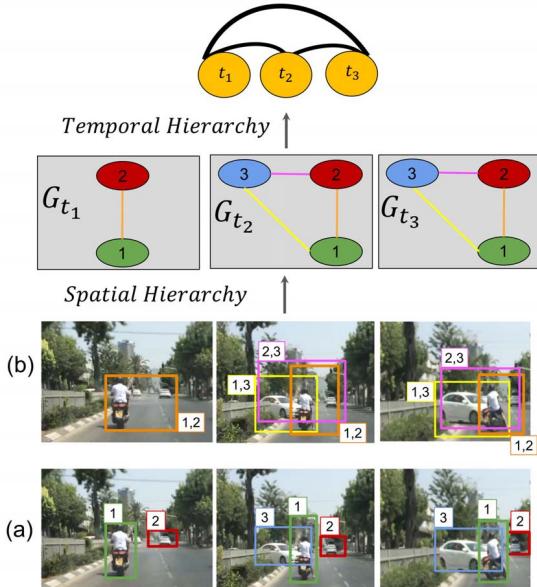


Figure 1: Our Spatio-Temporal Action Graph (STAG) approach models inter-object relations per frame. Each frame representation contains objects (a) and their relations (b). Our model contains two hierarchies: a spatial hierarchy to model all possible object interactions in a frame and a temporal hierarchy to aggregate the temporal context for the whole video.

may perform poorly in a test condition when observing a green pickup truck that is actually about to hit a red sedan. It is thus important to develop activity recognition models that can generalize effectively across object appearance and inter-object interactions.

Early deep learning approaches to activity recognition were limited to scene-level image representations, directly applying convolutional filters to full video frames, and thus not modeling objects or their interactions explicitly [7, 41, 37]. Although networks are growing deeper and wider, the method for extracting features from network backbones is often still a basic pooling operation, with or without a pixel-wise attention step. These conventional deep learning approaches are unable to directly attend to objects or their spatial relationships explicitly.

A natural approach to the above problem is to build mod-

*Equal Contribution.

†Work done during an internship at Nexar.

els that can capture relations between objects across time. This object-centric approach can decouple the object detection problem (for which more data is typically available, e.g., images of cars) and the problem of activity recognition. Many classic approaches to activity recognition explored object-based representations [12, 47, 26, 28, 4, 14]; yet with conventional learning methods such approaches did not show significant improvements in real-world evaluation settings. Several deep models have been recently introduced that directly represent objects in video in activity recognition tasks. For example [1] uses a relation network followed by an RNN, and [45] uses a spatio-temporal graph constructed whose nodes are detected objects. These models showed that a deep model with a dense graph defined over scene elements can lead to increased performance, but were limited in that only unary object appearance was considered, with a fully-connected spatio-temporal graph without taking object relations into account.

In this paper we propose a novel Spatio-Temporal Action Graph (STAG), which offers improved activity recognition. Our model design is motivated by the following two points. First, we observe that relations between objects are captured in the bounding box containing both objects more effectively than in the object boxes individually. Our graph utilizes explicit appearance terms for edges in the graph, forming a type of “visual phrase” term for each edge [32]: edge weights in our graph are formed using a descriptor pooled over the spatial extent of the union of boxes of each object pairs. Our experiments prove that modeling the visual appearance between objects outperforms other techniques (e.g. similarity and concatenation) for object-object interactions.

Second, the object interactions in one video are more concentrated in certain times which requires more structured spatial-temporal hierarchical feature representation. We propose a spatio-temporal disentangled feature embedding in our graph, factoring spatial and temporal connections into two hierarchies which first refine the edges considering all possible relations in a frame over space, and then over time. In our spatial hierarchy, the relations are refined by considering all possible relations within a frame and then aggregated to form per-frame descriptor. Next, we use the temporal hierarchy to aggregate the temporal context for the whole video and use it as input to the video classifier. We argue that this architecture is ideally structured for capturing relations that underlie typical actions in video. Indeed our empirical results show that it outperforms other spatio-temporal approaches without explicit hierarchy, including LSTM based ones.

Another key contribution of our work is a new dataset¹ for collision activity detection in driving scenario. Activ-

ity recognition is of key importance in the domain of autonomous driving, in particular detecting collisions or near-collisions is of utmost importance. Most of the research on this topic is in simulation mode, and we introduce the first attempt to studying it in real world data. Additionally, this is the first dataset containing object-object interactions, while the current activity recognition datasets [13, 35] mostly contain limited are human-object interactions that have small number of objects per scene. Thus, they cannot contain rich relation information. Here we provide a new dataset which will allow researchers to study recognition of such rare and complex events. Our Collision dataset was collected from real-world dashcam data consisting of 803 videos containing collisions or near-collisions from more than ten million rides.² In Fig. 1, we demonstrate our approach on a driving scene for collision event detection.

We evaluate our STAG model on both the Collision dataset as well the Charades [36] activity recognition benchmark, demonstrating improvement over previous baselines. Our results confirm that the use of explicit object representations in spatial-temporal hierarchy can offer better generalization performance for deep activity recognition in realistic conditions with limited training data.

2. Related Work

Video Activity Recognition. Early deep learning activity recognition systems were essentially “bag of words models”, where per-frame features are pooled over an entire video sequence [19]. Later work used recurrent models (e.g., LSTM), to temporally aggregate frame features [7, 50]. Another line of activity classifiers use 3D spatio-temporal filters to create hierarchical representations of the whole input video sequence [16, 40, 41, 42]. While spatio-temporal filters can benefit from large video datasets, they cannot be pretrained on images. Two-stream networks were proposed to leverage image pretraining in RGB as well as capturing fine low-level motion in another optical flow stream [37, 9]. I3D [5] was designed to inflate 2D kernels to 3D to learn spatio-temporal feature extractors from video while leveraging image pre-trained weights. In all these models, whole frame video features are extracted without using object and inter-object details as our model does.

Object interactions have been utilized for tackling various activity recognition tasks [23, 49, 39, 10, 27, 14], e.g. by using spatio-temporal tubes [12, 47], spatially-aware embeddings [26], and spatio-temporal graphical models [28, 4, 14, 11]. Probabilistic models have also been used in this context [31, 8, 48]. But with conventional learning methods the addition of explicit object detection models often did not show significant improvements in real-world evaluations. Recently, object interactions in adjacent frames were modeled in [25, 1] followed by RNNs for capturing

¹The publicly available dataset can be found at <https://github.com/roeiherz/STAG-Nets>.

²There are only relatively few collision videos, since naturally, such events are rare.

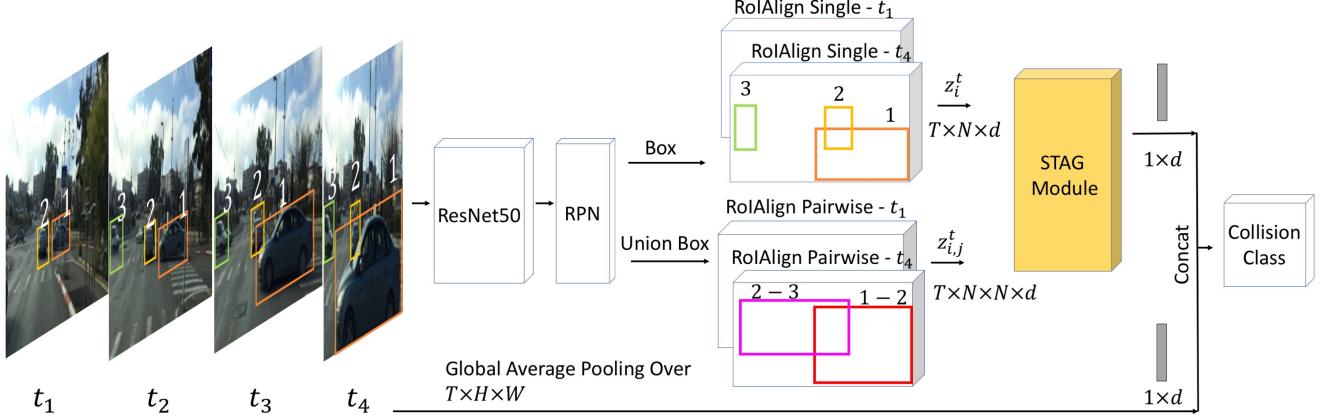


Figure 2: STAG Network Architecture. Our STAG network architecture is comprised of: (1) A backbone of ResNet50 and RPN produces a set of bounding boxes proposals for each frame. (2) An *RoIAlign Single* layer extracts z_i^t features from the backbone for the boxes. In parallel, every pair of box proposals is used for computing a union box, and pairwise features $z_{i,j}^t$ are extracted similar to the z_i^t features. (3) The z_i^t and $z_{i,j}^t$ are used as inputs to an Spatio-Temporal Action Graph Module (see Fig. 3) which outputs a d dimensional feature representing the entire sequence.

temporal structure. Also, Wang [46] proposes to represent videos as space-time region graphs and perform reasoning on this graph representation via Graph Convolutional Networks. In [46] objects in one video are allowed to interact with each other without constraints while we enforce more structured spatial-temporal feature hierarchy for better video feature encoding.

Graph Neural Networks and Self-attention. Recently graph neural networks have been successfully applied in many computer vision applications: visual relational reasoning [2, 51, 1, 29, 22], image generation [17] and robotics [33]. Message passing algorithms have been redefined as various graph convolution operations [15, 21]. The graph convolutional operation is essentially equivalent to “non-local operation” [44] derived from the self-attention concept [43]. In this paper, we take “non-local operation” as our graph convolutional operation.

Autonomous Driving. Deep learning has recently been applied to learn autonomous-driving policies [6, 3]. Collision avoidance is an important goal for self-driving systems [18]. Collision vision data is difficult to collect in the real world since these are unexpected rare events. [20] tackles the collision data scarcity by simulation. However synthetic data is still very different from real data, and hence training on simulation is not always sufficient. In this paper, we introduce a challenging collision dataset based on real-world dashcam data. At the same time, we propose a model suitable for classifying rare events by modeling the key object interactions with limited training examples.

3. Spatio-Temporal Action Graphs

In this section we describe our proposed Spatio-Temporal Action Graph Network (STAG). The overall ar-

chitecture is shown in Fig. 2 and the STAG module is further described in Fig. 3.

We begin with some definitions. The following constants are used: T is the number of frames, N is the maximum number of objects (i.e., bounding boxes) per frame, and d is the feature dimensionality (i.e., the dimension of bounding boxes descriptors). We also use $[T]$ to denote the set of input frames $\{1, \dots, T\}$ and $[B]$ to denote the set of bounding boxes in each frame $\{1, \dots, N\}$. At a high level the model proceeds in the following stages:

Detection Stage - The image is pre-processed with a detector to obtain features for each bounding box (i.e., objects) and each pair of boxes (i.e., relations).

Spatial Context Hierarchy Stage - Each relation feature is refined using context from other relations, and all relation features are summarized in a single feature per frame.

Temporal Context Hierarchy Stage - Each frame feature is refined using context from other frames, and all frame features are summarized in a single feature, which is then used for classification.

3.1. Detection Stage

Before applying the two disentangled spatial and temporal context aggregations, we construct one initial graph representation encoding objects and their relations. As a first step, we detect region proposal boxes and extract corresponding box features through an RoIAlign layer of a Faster R-CNN [30]. Instead of only using object features with their bounding boxes, we believe that the spatial relation of each pair of bounding box should be important and encoded into the initial graph representation for subsequent spatial and temporal context aggregation. Specifically, for each pair of boxes, we consider its union (see Fig. 2) and use an RoIAlign layer to extract the union box features as

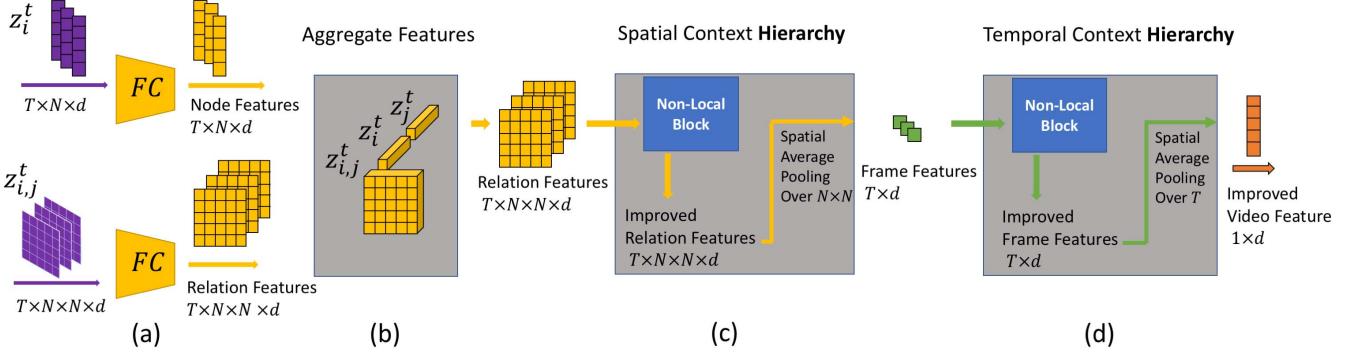


Figure 3: STAG Module: (a) First, the initial node features z_i^t and union box relation features $z_{i,j}^t$ are embedded into a feature space with dimension d . (b) The pair of node features and their relation features are concatenated to create an aggregated representation with objects and their interactions. These are then embedded into dimension d using a fully-connected network. (c) The **Spatial Context Hierarchy Stage** uses a Non-Local block to obtain improved relation features and then pools over the spatial dimension to obtain per-frame features. (d) Then, the **Temporal Context Hierarchy Stage** uses a Non-Local block to calculate improved frame features and pools over the temporal dimension. The final output is a video-level feature containing the entire scene information and video dynamics.

initial relation features with the union boxes capturing the spatial appearance of each pair of objects. This results in two sets of tensors shown in Fig. 3a:

- Single-object features: For each time step $t \in [T]$ and box $i \in [B]$ we have a feature vector $z_i^t \in \mathbb{R}^d$ for the corresponding box. The feature contains the output of the RoIAlign layer.
- Object-pair features: For each time step $t \in [T]$ and box-pair $i \in [B], j \in [B]$ we have a feature vector $z_{i,j}^t \in \mathbb{R}^d$ for the corresponding pair of boxes. The feature contains the output of the RoIAlign layer.

We thus have a tensor of size $T \times N \times d$ for single object features, and a tensor of size $T \times N \times N \times d$ for object-pair features. Then (Fig. 3b) we concatenate each relation (object-pair) feature with the two corresponding node features and embed the result into dimension d (using an FC layer) to form one aggregated representation with objects and their interactions. The resulting tensor is thus of size $T \times N \times N \times d$.

3.2. The STAG Module

The output of the Detection Stage is a tensor of size $T \times N \times N \times d$ with features for each object interaction in each frame. In what follows, we describe how these are refined and reduced in the two complementary hierarchies of space and time. Before we introduce these stages, we first recap the non-local operation from [44]. These are an efficient, simple and generic component for capturing long range dependencies. Formally, given a set \mathcal{V} of vectors v_1, \dots, v_k , the non-local operator transforms these into a new set \mathcal{V}' of vectors v'_1, \dots, v'_k via the function:

$$v'_i = \frac{1}{C(\mathcal{V})} \sum_{\forall j} f(v_i, v_j)g(v_j) \quad (1)$$

where $C(\mathcal{V})$ is a normalization factor and f and g are learned pairwise and singleton functions. We next describe the final two stages of the STAG model.

Spatial Context Hierarchy Stage. The goal of this stage is twofold. First, it refines the relation features so that each feature incorporates information from all the other relations. This is done by applying a non-local operation to all the N^2 feature vectors (each of dimension d) that are the output of the Detection Stage. The outcome is another tensor of size $T \times N \times N \times d$. Next, it generates a single feature representing the relation information in the frame, by average pooling the above tensor, resulting in a tensor of size $T \times d$. See Fig. 3c. We visualize some of the object proposals and their relations in Fig. 5.

Temporal Context Hierarchy Stage. At this stage, information from all frames is integrated into a single vector. This is done by applying a Non-Local block to the T vectors (each of dimension d) that are the output of the Spatial Context Hierarchy Stage. The final output is a single d dimensional feature vector capturing whole video information obtained by average pooling the above tensor. See Fig. 3d. We visualize some of the frames and their relations in Fig. 6.

4. The Collision Dataset

We introduce the Collision dataset comprised of real-world driving videos. Such a dataset is valuable for developing autonomous driving models. Using videos, and specifically visual information, is important for accurate and timely prediction of collisions. The dataset contains rare collision events from diverse driving scenes, including urban and highway areas in several large cities in the US. These events encompass collision scenarios (i.e., scenarios involving the contact of the dashcam vehicle with a fixed or moving object) and the near-collision scenarios (i.e., scenar-



Figure 4: Example visualization from the collision dataset: (a) extreme weather conditions such as snow and heavy raining. (b) near-collision with a truck/a bicycle rider. (c) day and night collisions.

Party type	dist.
Vehicle	85%
Bike	6%
Pedestrian	6%
Road object	1%
Motorcycle	1%

Weather	dist.
Clear	93%
Rain	5.3%
Snow	1.7%

Lighting	dist.
Day	62%
Night	38%

Table 1: Collision dataset statistics: involved party, weather, and lighting conditions.

ios requiring an evasive maneuver to avoid a crash). Such driving scenarios most often contain interactions between two vehicles, or between a vehicle and a bike or pedestrian. Classifying such events therefore naturally requires modeling object interactions, which was our motivation for developing the STAG model. We will release a publicly available challenge based on this dataset upon acceptance of paper.

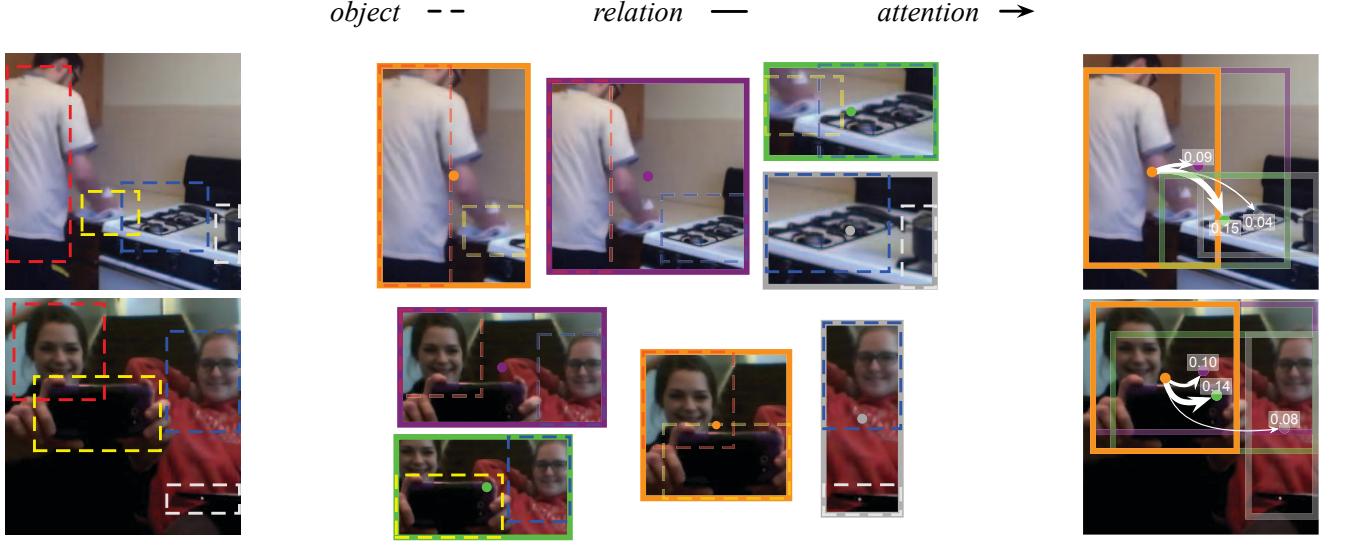
Data collection. The data was collected from a large-scale deployment of connected dashcams. Each vehicle is equipped with a dashcam and a companion smartphone app that continuously captures and uploads sensor data such as IMU and gyroscope readings. Overall, the vehicles collected more than 10 million rides, and rare collision events were automatically detected using a triggering algorithm based on the IMU and gyroscope sensor data.³ The events were then manually validated by human annotators based on visual inspection to identify edge case events of collisions and near-collisions, as well as non-risky driving events. Of all the detected triggers, our subset contains 743 collisions

and 60 near-collisions from different drivers. Each video clip contains one such event typically occurring in the middle of the video clip. The clip duration is approximately 40 seconds on average and the frame resolution is 1280×720 .

The full and few-shot datasets. We created two data versions. The full dataset contains a total of 803 videos with 732 videos as training data (44 of them are near collision) and 71 videos as test data (6 of them are near-collision). We use a relatively low frame rate of 5 fps to convert video clips to frames, in order to avoid using near-duplicate frames. Each clip is broken into three segments to train our model. Specifically, we split each video into three segments of 20 frames each: two negative segments (non-risky driving scenes) and one positive segment (a collision scene) for each collision event. The positive segment is sampled at the time of collision, and the two non-overlapping negative segments are sampled before the time of collision, since after collision the scene is already in a collision state. For the near-collision event, we sample three negatives since there is no positive segment in the video clip. After this processing, we have a total 2409 video segments, out of which 1656 are negative examples and 753 are positives. The few-shot dataset is purposely designed to motivate the development of few-shot recognition algorithms. It contains 125 videos, with 25 training videos and 100 testing videos. Data processing is the same as the full-version. The positive to negative ratio in both versions of dataset is approximately 1 : 3.

³The algorithm is tuned to capture driving maneuvers such harsh braking, acceleration, and sharp cornering.

Diversity. Our dataset is collected for recognizing collisions in natural driving scenes. To get an intuitive feeling



(a) Object Proposals

(b) Relations

(c) Relation Attentions

Figure 5: Spatial context hierarchy processing. Our model first extracts object proposals (see Fig. 3b) and then composes them into relations by taking the union of each possible object-pair proposal pair. Our model then infers the per-frame relation interactions via an attention step. (c) illustrates the attention on the top-3 relation-to-relation interactions and their scores for the relation box colored by orange.

for our collision dataset, we visualize several video examples in Fig. 4. The coverage of the dataset includes various types of weather conditions (Fig. 4a), the parties involved (Fig. 4b) and lighting conditions (Fig. 4c). We use visual inspection to analyze the identity of parties involved in the collisions in Tab. 1. We find that most of the data (85%) consists of crashes involving two vehicles, and collisions involving pedestrians and cyclists takes up to 6% each. Tab. 1 also shows the distribution of weather conditions and lighting conditions. With a majority of clear weather (93%), the extreme rain and snow video clips take 5.3% and 1.7% each. Day-time takes 62% with the rest of 38% being night-time.

5. Experiments on the Collision Dataset

Our method is designed to address rich inter-object interactions. The only dataset that captures these as Charades, whereas other datasets contain limited human-object interactions. It was thus natural to evaluate it on Charades and Collision. We next describe the application of our STAG model to those datasets.

5.1. Implementation Details

Model Details. We use Faster R-CNN with ResNet50 as a backbone, taking a sequence of $T = 20$ frames, and generating bounding box proposals for each of the T frames. Specifically, the strides in FPN are set as the same as [24].

The input frames are resized to the maximum dimension of 256 with padding. Considering the training time

and memory limit, we take the top $N = 12$ region proposals on each frame after non-maximum suppression with IoU threshold 0.7, which are sufficient for capturing the semantic information on the Collision dataset. Features for the $N = 12$ objects and $N \cdot N$ object interaction relations are extracted following Sec. 3, resulting in feature representations z_i for objects and $z_{i,j}$ for relations.

Training and Inference. We train STAG using SGD with momentum 0.9 and an initial learning rate 0.01. The learning rate is decayed by a factor of 0.5 each epoch, and gradient is clipped at norm 5. Each batch includes a video segment of T frames. Two kinds of ground truth data are utilized during training: the ground truth bounding box annotations on each frame and the collision label per segment.

The loss for the STAG model contains two components: the bounding box localization related losses used in the Faster-RCNN detector and the multi-class activity classification loss, as is standard with two stage detectors. To train our STAG model for collision prediction, we apply a binary cross entropy loss between the binary collision prediction logit and the ground-truth collision label.

5.2. Model Variants

The STAG model progressively processes the box features and the spatial appearance features of pairwise boxes $z_i^t, z_{i,j}^t$ into a single vector for final activity recognition. To explore the importance of the spatial and temporal aspects of STAG, we consider the following variants:

- (1) LSTM Spatial Graph - We study the effect of the

	Accuracy	
	Full Dataset	Few-shot Dataset
I3D	82.4	76
C3D	79.9	72
LSTM Spatial Graph	77.5	67
LSTM boxes	69.5	69
STAG	84.5	76.3

Table 2: Classification accuracy on the Collisions dataset for the STAG model and its variants, and the C3D & I3D model.

	Accuracy	
	Full Dataset	Few-shot Dataset
LSTM spatial Graph	83.56	73.1
LSTM boxes	81.2	72.3
STAG	85.5	76.7

Table 3: Classification accuracy on the Collisions dataset for STAG model and variants, when averaged with the C3D model.

STAG “Temporal Context Hierarchy” stage, as compared to a recurrent neural network based solution. To do so, we replace the “Temporal Context Hierarchy” stage with an LSTM that processes the same tensor of size $T \times d$.

(2) **LSTM Boxes** - We study the effect of the “Spatial Context Hierarchy” stage by replacing it with average pooling of the node features, to obtain a tensor of size $T \times d$. We also train two other popular activity recognition models on Collision dataset: the C3D model [41] and I3D model [5]. We used pretrained weights for C3D and I3D. The C3D was pretrained on Sports-1M while the I3D was pretrained on Kinetics.

5.3. Results

We first compare the STAG results on the full dataset to the model variants described in Sec. 5.2. Tab. 2 reports classification accuracy. Firstly, STAG outperforms all the other models including C3D and I3D. Replacing the temporal processing in STAG with an LSTM as in LSTM SPATIAL GRAPH, we get 7% accuracy decrease, showing the superiority of our temporal modeling over LSTM. Further removing the pairwise object modeling, we see accuracy further decrease by 8% in LSTM BOXES.

Finally, we consider a simple ensemble model of STAG and C3D by simply averaging their output scores. Results of this combination are shown in Tab. 3. We can see the combination improves the original C3D accuracy, showing the benefits of object interaction modeling. Among all the ensemble results, the STAG model still maintains the highest accuracy result 85.5%.

We also show the results on the few-shot dataset in Tab. 2 and Tab. 3. It can be seen that the two LSTM model variants almost fail on this challenging dataset. Although our STAG model achieves marginal improvement compared to

	Edge	Hierarchy	Accuracy
STAG Cat	Node concat.	Space & Time	83.1
STAG Sim	Cosine sim.	Space & Time	83.5
STAG Time	Node interactions	Time only	78.8
STAG Space	Node interactions	Space only	82.6
STAG	Node interactions	Space & Time	84.5

Table 4: Hierarchy & Edge features Ablations on the Collision dataset. “Node Interactions” refers to using relations features for the edge features.

the C3D and I3D, the relative low accuracy numbers highlight the challenges of this setting.⁴ We encourage the community to further develop few-shot based activity recognition models to tackle this challenging few-shot dataset.

5.4. Ablation Studies

We also design some direct ablation studies for the components in our STAG model. To validate the effectiveness of our disentangled spatio-temporal hierarchies, we design two ablation studies for the two attention hierarchies: (1) STAG Space - Replacing the spatial hierarchy by directly pooling. (2) STAG Time - Replacing the temporal hierarchy by directly pooling.

The results are shown in Tab. 4. It can be seen that both ablations decrease accuracy, but that the temporal hierarchy has a larger effect on performance.

In addition to our visual appearance relation features, we explore the use of different relation features in Tab. 4: (1) STAG Cat - Set edge feature to be just the concatenation of the corresponding node features (i.e., union box is not used). (2) STAG Sim - Set edge feature to be cosine similarity of the two corresponding node features (see [46]).

Both methods result in approximately one point accuracy decrease, indicating the superiority of using spatial appearance features of union boxes as edge features in our hierarchical STAG models.

6. Experiments on the Charades Dataset

To further validate the effectiveness of our model on publicly available action recognition benchmarks, we also evaluate it on the Charades dataset [36]. We follow the official split (8K training and 1.8K validation videos) to train and test our model. The average video duration is around 30 seconds with 157 multiple action classes and we report our results by the metric of mean Average Precision (mAP).

We follow the same experiment setup as described in STRG (Spatio-Temporal Region Graph) [46] and use a backbone network of ResNet-50 Inflated 3D ConvNet (I3D) [44] for all of our experiments.

⁴We note however, that all these models are not specifically designed for the few-shot setting.



Figure 6: Our Temporal Context Hierarchy. After aggregating spatial features to represent each frame separately (see Fig. 3c), our model learns relations across the frames (see Fig. 3d). In each row, we show eight consecutive frames in a video in which two frames are highlighted in red or blue. We use arrows to represent the temporal attention across frames which started from either of highlighted frames. The attention score associated with them are noted beside arrows.

Training and Inference. Our network takes 32 video frames as inputs which are sampled at 6fps, resulting in maximum input duration of about 5 seconds. We use a spatial resolution of 224×224 . Data augmentation is as in [38]. The top $N = 15$ object proposals are selected.

To train our model, we follow the same training schedule as specified in STRG using a mini-batch of 8 videos for each iteration and repeat it with 100K iterations in total. The training objective is a simple cross entropy loss. During inference, we apply multi-crop testing [44, 46] for better performance and the final recognition results are based on late fusion of classification scores.

Results. Tab. 5 compares STAG to various baselines on Charades. It can be seen that it compares favorably with prior works that used the same ResNet-50 I3D backbone: STAG improves 5.4% over the I3D model, and 1.0% over STRG.

Ablations. Next, we run the same ablation studies (STAG Space, STAG Time, STAG Cat) as in Sec. 5.4. Results are shown in Tab. 6. It can be seen that as with the Collisions data, all STAG ablations result in decreased accuracy. STRG Sim refers to the STRG model which uses cosine similarity between the nodes as the edge features, while either not discriminating the nodes from different frames at all or only applying heuristic backward-forward node association as space-time hierarchy. We compare STRG Sim to STAG Relation, a model that uses the same relation features as STRG, and thus it is a direct comparison between the similarity from [46] and our relation features approach. Our design of relations feature as the edge feature captures object interactions and brings a 0.6% performance gain over STRG Sim.

7. Conclusion

The interaction of objects over time is often a critical cue for understanding activity in videos. We presented a

	Backbone	Modality	mAP
2-Steam [37]	VGG-16	RGB w/ Flow	18.6
2-Steam w/ LSTM [37]	VGG-16	RGB w/ Flow	18.6
Async-TF [34]	VGG-16	RGB w/ Flow	22.4
a Multiscale TRN [52]	Inception	RGB	32.9
I3D [5]	Inception	RGB	32.9
I3D [46]	R50-I3D	RGB	31.8
STRG [46]	R50-I3D	RGB	36.2
STAG (ours) [44]	R50-I3D	RGB	37.2

Table 5: Classification mAP in the Charades dataset. [36]

	Edge	Hierarchy	mAP
I3D	-	-	31.8
STRG Sim [46]	Cosine sim.	No hierarchy	35.0
STRG [46]	Cosine sim.	Space-Time Heuristic	36.2
STAG Relation	Node interactions	No hierarchy	35.6
STAG Cat	Node concat.	No hierarchy	34.5
STAG Space	Node interactions	Space only	34.7
STAG Time	Node interactions	Time only	36.6
STAG	Node interactions	Space & Time	37.2

Table 6: Hierarchy & Edge features Ablations on the Charades dataset.

novel inter-object graph representation which included explicit appearance models for edge-terms in the graph as well as a novel factored embedding of the graph structure into spatial and temporal representation hierarchies. We demonstrated the effectiveness of our model on the Charades activity recognition dataset as well as on a new dataset of driving near-collision events; our model significantly improved performance compared to baseline approaches without object-graph representations or with previous graph-based models.

Acknowledgements

This work was completed in partial fulfillment for the Ph.D degree of the first author.

References

- [1] F. Baradel, N. Neverova, C. Wolf, J. Mille, and G. Mori. Object level visual reasoning in videos. In *European Conf. Comput. Vision*, pages 105–121, 2018. [2](#), [3](#)
- [2] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. [3](#)
- [3] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016. [3](#)
- [4] W. Brendel and S. Todorovic. Learning spatiotemporal graphs of human activities. In *Computer vision (ICCV), 2011 IEEE international conference on*, pages 778–785. IEEE, 2011. [2](#)
- [5] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733, 2017. [2](#), [7](#), [8](#)
- [6] C. Chen, A. Seff, A. Kornhauser, and J. Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2722–2730, 2015. [3](#)
- [7] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015. [1](#), [2](#)
- [8] A. Fathi, A. Farhadi, and J. M. Rehg. Understanding egocentric activities. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 407–414. IEEE, 2011. [2](#)
- [9] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1933–1941, 2016. [2](#)
- [10] R. Girdhar, J. Carreira, C. Doersch, and A. Zisserman. Video action transformer network. *arXiv preprint arXiv:1812.02707*, 2018. [2](#)
- [11] R. Girdhar, J. a. Carreira, C. Doersch, and A. Zisserman. Video Action Transformer Network. In *CVPR*, 2019. [2](#)
- [12] G. Gkioxari and J. Malik. Finding action tubes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 759–768, 2015. [2](#)
- [13] R. Goyal, S. E. Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag, et al. The “something something” video database for learning and evaluating visual common sense. In *ICCV*, volume 1, page 3, 2017. [2](#)
- [14] A. Gupta, A. Kembhavi, and L. S. Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(10):1775–1789, 2009. [2](#)
- [15] R. Herzig, M. Raboh, G. Chechik, J. Berant, and A. Globerson. Mapping images to scene graphs with permutation-invariant structured prediction. *arXiv preprint arXiv:1802.05451*, 2018. [3](#)
- [16] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):221–231, 2013. [2](#)
- [17] J. Johnson, A. Gupta, and L. Fei-Fei. Image generation from scene graphs. *arXiv preprint arXiv:1804.01622*, 2018. [3](#)
- [18] G. Kahn, A. Villaflor, V. Pong, P. Abbeel, and S. Levine. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017. [3](#)
- [19] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. [2](#)
- [20] H. Kim, C. Suh, K. Lee, and G. Hwang. Crash to not crash: Learn to identify dangerous vehicles using a simulator. In *Association for the Advancement of Artificial Intelligence (AAAI). THE ASSOCIATION FOR THE ADVANCEMENT OF ARTIFICIAL INTELLIGENCE*, 2019. [3](#)
- [21] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. [3](#)
- [22] R. Krishna, I. Chami, M. S. Bernstein, and L. Fei-Fei. Referring relationships. *ECCV*, 2018. [3](#)
- [23] L.-J. Li and L. Fei-Fei. What, where and who? classifying events by scene and object recognition. 2007. [2](#)
- [24] T. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 936–944, 2017. [6](#)
- [25] C.-Y. Ma, A. Kadav, I. Melvin, Z. Kira, G. AlRegib, and H. P. Graf. Attend and interact: Higher-order object interactions for video understanding. *arXiv preprint arXiv:1711.06330*, 2017. [2](#)
- [26] P. Mettes and C. G. Snoek. Spatial-aware object embeddings for zero-shot localization and classification of actions. In *Proceedings of the International Conference on Computer Vision*, 2017. [2](#)
- [27] D. J. Moore, I. A. Essa, and M. H. Hayes. Exploiting human actions and object context for recognition tasks. Technical report, Georgia Institute of Technology, 1999. [2](#)
- [28] B. Packer, K. Saenko, and D. Koller. A combined pose, object, and feature model for action understanding. In *CVPR*, pages 1378–1385. Citeseer, 2012. [2](#)
- [29] M. Raboh, R. Herzig, G. Chechik, J. Berant, and A. Globerson. Learning latent scene-graph representations for referring relationships. *ArXiv e-prints*, 2019. [3](#)
- [30] S. Ren, K. He, R. B. Girshick, and J. Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *CoRR*, abs/1506.01497, 2015. [3](#)
- [31] M. S. Ryoo and J. K. Aggarwal. Hierarchical recognition of human activities interacting with objects. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [2](#)

- [32] M. A. Sadeghi and A. Farhadi. Recognition using visual phrases. 2011. [2](#)
- [33] A. Sanchez-Gonzalez, N. Heess, J. T. Springenberg, J. Merel, M. Riedmiller, R. Hadsell, and P. Battaglia. Graph networks as learnable physics engines for inference and control. *arXiv preprint arXiv:1806.01242*, 2018. [3](#)
- [34] G. A. Sigurdsson, S. Divvala, A. Farhadi, and A. Gupta. Asynchronous temporal fields for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 585–594, 2017. [8](#)
- [35] G. A. Sigurdsson, A. Gupta, C. Schmid, A. Farhadi, and K. Alahari. Charades-ego: A large-scale dataset of paired third and first person videos. *arXiv preprint arXiv:1804.09626*, 2018. [2](#)
- [36] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, pages 510–526. Springer, 2016. [2](#), [7](#), [8](#)
- [37] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014. [1](#), [2](#), [8](#)
- [38] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [8](#)
- [39] C. Sun, A. Shrivastava, C. Vondrick, K. Murphy, R. Sukthankar, and C. Schmid. Actor-centric relation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018. [2](#)
- [40] G. W. Taylor, R. Fergus, Y. LeCun, and C. Bregler. Convolutional learning of spatio-temporal features. In *European conference on computer vision*, pages 140–153. Springer, 2010. [2](#)
- [41] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. [1](#), [2](#), [7](#)
- [42] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1510–1517, 2018. [2](#)
- [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. [3](#)
- [44] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7794–7803, 2018. [3](#), [4](#), [7](#), [8](#)
- [45] X. Wang and A. Gupta. Videos as space-time region graphs. *arXiv preprint arXiv:1806.01810*, 2018. [2](#)
- [46] X. Wang and A. Gupta. Videos as space-time region graphs. In *ECCV*, 2018. [3](#), [7](#), [8](#)
- [47] P. Weinzaepfel, X. Martin, and C. Schmid. Towards weaklysupervised action localization. *arXiv preprint arXiv:1605.05197*, 3(7), 2016. [2](#)
- [48] J. Wu, A. Osuntogun, T. Choudhury, M. Philipose, and J. M. Rehg. A scalable approach to activity recognition based on object use. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8. IEEE, 2007. [2](#)
- [49] B. Yao and L. Fei-Fei. Recognizing human-object interactions in still images by modeling the mutual context of objects and human poses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(9):1691–1703, 2012. [2](#)
- [50] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015. [2](#)
- [51] V. Zambaldi, D. Raposo, A. Santoro, V. Bapst, Y. Li, I. Babuschkin, K. Tuyls, D. Reichert, T. Lillicrap, E. Lockhart, et al. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*, 2018. [3](#)
- [52] B. Zhou, A. Andonian, A. Oliva, and A. Torralba. Temporal relational reasoning in videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 803–818, 2018. [8](#)