

Unified Graph Structured Models for Video Understanding

Anurag Arnab Chen Sun Cordelia Schmid
Google Research

{aarnab, chensun, cordelias}@google.com

Abstract

Accurate video understanding involves reasoning about the relationships between actors, objects and their environment, often over long temporal intervals. In this paper, we propose a message passing graph neural network that explicitly models these spatio-temporal relations and can use explicit representations of objects, when supervision is available, and implicit representations otherwise. Our formulation generalises previous structured models for video understanding, and allows us to study how different design choices in graph structure and representation affect the model’s performance. We demonstrate our method on two different tasks requiring relational reasoning in videos – spatio-temporal action detection on AVA and UCF101-24, and video scene graph classification on the recent Action Genome dataset – and achieve state-of-the-art results on all three datasets. Furthermore, we show quantitatively and qualitatively how our method is able to more effectively model relationships between relevant entities in the scene.

1. Introduction

Deep learning has enabled rapid advances in many image understanding tasks such as image classification [20], object detection [45] and semantic segmentation [7]. However, progress on recent video understanding datasets such as AVA [17] and Charades [48] has lagged behind in comparison. Further progress in the video understanding tasks posed by these datasets would facilitate applications in autonomous vehicles, health monitoring and automated media analysis and production among others.

A reason why video understanding is so challenging is because, as shown in Fig. 1, it requires understanding the interactions between actors, objects and other context in the scene. Furthermore, these interactions are not always observable from a single frame, and thus require reasoning over long temporal intervals. This is illustrated in Fig. 1, where understanding the actions of the person in the centre-frame is not possible from the target keyframe alone. In order to know that the woman is “listening”, we need to con-

sider the man who is speaking but no longer in the scene. And to correctly infer that the woman is “driving” the car, rather than “riding” like the man, we must take into account that she is later holding the steering wheel.

Video is a significantly higher-dimensional signal than single images, due to its additional temporal axis, and so we believe learning these unlabelled interactions directly from current datasets with large convolutional networks is not feasible. In this paper, we propose a structured graph neural network to explicitly model these spatio-temporal interactions. We model actors and objects (explicitly with bounding boxes when we have supervision, and implicitly otherwise) as nodes in our spatio-temporal graph and perform message passing inference to directly model their relations.

Although a wide range of graph-structured models have been proposed for action recognition, we note that there has been no unifying formulation for these models. As such, some works only model spatial relations between actors and objects [14, 53], but not how these interactions evolve over time. Other approaches model long-range temporal interactions [60], but do not capture spatial relations and are not trained end-to-end. And whilst some methods do model spatio-temporal interactions of objects [3, 59], their explicit representations of objects need additional supervision, and are not evaluated on spatio-temporal localisation tasks which requires detailed understanding and is necessary for analysing untrimmed videos.

Our graph network formulation based on the message-passing neural network [13] abstraction, allows us to explicitly model interactions between actors, objects and the scene, and how these interactions evolve over time. Our flexible model allows us to use explicit object representations from a pretrained Region Proposal Network (RPN) [45], and/or implicitly from convolutional feature maps without additional supervision. Moreover, our general formulation allows us to interpret previous work [14, 53, 59, 60, 69] as special cases, and thus understand how different design choices in object representation, graph connectivity and message passing functions affect the model’s performance. We demonstrate our versatile model on two different tasks: spatio-temporal action detection on AVA [17]



Figure 1: Understanding videos requires reasoning over long-term spatio-temporal interactions between actors, objects and the environment. The actions of the woman in the centre frame are ambiguous given the nearby frames which typical 3D CNN architectures consider. However, by considering her interactions with the man from nearby frames, we know she is “listening” to him. And the fact that she is later holding the steering wheel indicates she is “driving” in contrast to the man who is “riding”. In this paper, we propose a spatio-temporal graph on which we perform message passing to explicitly model these spatio-temporal interactions. Example from the AVA dataset [17].

and UCF101-24 [52], and video scene graph prediction on the recent Action Genome [23] dataset. Both of these tasks require modelling the spatio-temporal interactions between the actors and/or objects in the scene, and we show consistent improvements from using each component of our model and achieve state-of-the-art results on each dataset. Furthermore, we observe that the largest improvements in AVA are indeed achieved on action classes involving human-to-human and human-to-object interactions, and visualisations of our network show that it is focusing on scene context that is intuitively relevant to its action classification.

2. Related Work

Modelling contextual relationships has a long history in scene understanding. Relevant examples of early works in this area included modelling interactions between humans and objects [18, 66], different objects [44] and relationships between human actions and scene context [40, 16]. Furthermore, it has also been shown that human vision is reliant on context too [42]. In this paper, we consider video understanding tasks, specifically spatio-temporal action recognition and video scene graph parsing, which involve reasoning about interactions between actors, objects and their environment in both space and time.

Early work in action recognition used hand-crafted features to encode motion information [32, 57]. Advances in deep learning first saw repurposing of 2D image convolutional neural networks (CNNs) for video as “two stream” networks [27, 49] followed by spatio-temporal 3D CNNs [6, 12, 54, 70]. However, these architectures focus on extracting coarse, video-level features and are not suitable for learning the fine-grained relations depicted in Fig. 1. Consequently, whilst initial approaches to spatio-temporal action detection involved extending 2D object detectors [37, 45] temporally [26, 50, 1, 63], current leading methods [14, 60, 69] on the AVA dataset [17] all explicitly model relationships, with approaches which we show can be interpreted as variants of graph neural networks.

Graph neural networks (GNNs) explicitly model relations between entities by modelling them as nodes in a

directed or undirected graph [5, 47, 29], which interact via a neighbourhood defined for each node. The self-attention [55] and Non-local [58] operators can also be thought of as GNNs [5, 13] where each element in a feature map is a node, and all nodes are fully-connected to each other. Such attention-based models have excelled at a number of natural language processing and computer vision tasks and have inspired many follow-up methods [21, 8, 9, 61, 68].

Many structured models have also recently been employed in video understanding. However, there has previously been no coherent framework to unify these approaches. Consequently, some works only model spatial relations between actors and objects, but not how these evolve over time [53, 14]. And whilst LFB [60] models long range temporal interactions, it does not capture spatial relationships within a keyframe. Moreover, in order to model long-range interactions, [60] uses precomputed features and is thus not trained end-to-end. Our proposed method, which is based on the message passing neural network (MPNN) [13] framework, coherently models both spatial and temporal interactions. After describing our model in Sec. 3, we show how previous structured models for action detection [53, 14, 59, 60, 69] can be thought of as special cases of our model. Furthermore, our flexible model can reason about objects both when we have explicit supervision for them, and when we do not. Prior work in video understanding, which have proposed less generic models, have either assumed the case of explicit object supervision [3, 59] or not [53, 14, 69], but have not considered the scenario when both options are available. Furthermore, our general formulation allows us to ablate graph modelling design choices, such as the object representation, message passing functions and temporal context, in a manner not possible with these previous, more specific approaches. Moreover, we note that some approaches which have employed spatio-temporal graphs in video have only considered frame-to-frame interactions [22, 59, 3, 39, 41, 43], rather than long-range relations, and have not evaluated on spatio-temporal localisation which requires more detailed understanding

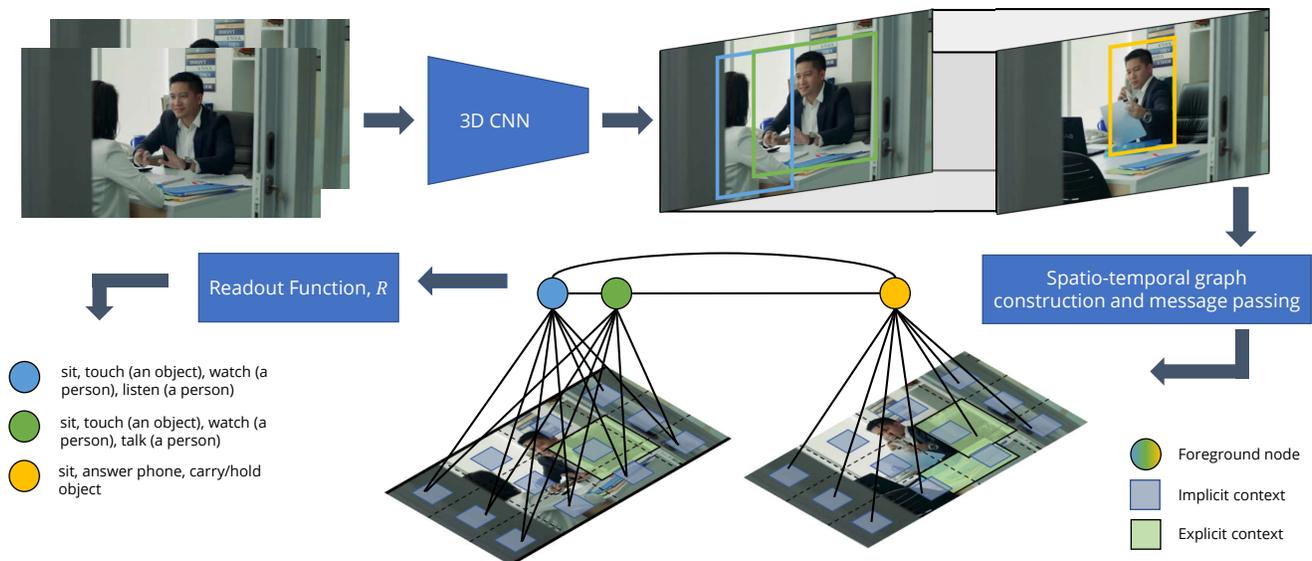


Figure 2: Overview of our method: We construct a spatio-temporal graph, and perform message-passing inference on it, to model interactions between actors, objects and their environment. Foreground nodes (circles) have readout functions associated with them for the task of interest (*i.e.* for action recognition, the nodes represent person bounding boxes which are classified into actions). Context nodes (squares) model additional information, and can either be implicit, as cells of the original feature map, or explicit by ROI-pooling external region proposals (RPN [45] not shown for clarity). The initial state of each node is a spatio-temporal feature vector extracted from a 3D CNN.

and is essential for analysing untrimmed videos.

We also note that scene graph parsing [25, 30] is another task that evaluates a model’s ability to reason about the interactions between different objects by representing objects as nodes and relationships as edges in a graph. Although the task was originally posed for single images [30], the recent Action Genome [23] dataset extends this task to video by adding annotations to Charades [48]. While GNN-based approaches have also been used in scene graph parsing [10, 33, 62, 64] for single images, to our knowledge, they have not been applied to model temporal relations in video. Moreover, the same model has not also been demonstrated on spatio-temporal action recognition like our method.

3. Proposed Approach

Our model aims to build a structured representation of a video by representing it as a graph of actors, objects and contextual elements in the scene, as shown in Fig. 2. This structured representation is then used to perform tasks which require understanding the interactions between the elements in the graph, such as action recognition and scene graph prediction. Note that we do not assume we have annotations for relevant scene context.

Our approach is based on Message Passing Neural Networks (MPNN) [13], as it is a flexible framework that generalises many previous graph neural network algorithms [4, 5, 29, 56] We review this approach, and describe how we adapt it for video understanding in Sec. 3.1, before detailing our model in Sec. 3.2 through 3.5. Finally, we dis-

cuss how previous structured models for video understanding [14, 53, 59, 60, 69] can be regarded as specific instantiations of our model in Sec. 3.6.

3.1. Message Passing Neural Networks (MPNNs)

MPNNs operate on a directed or undirected graph, \mathcal{G} consisting of nodes, $v \in \mathcal{V}$, and a neighbourhood for each node, \mathcal{N}_v , that defines the graph’s connectivity. For video models we distinguish between the spatial, \mathcal{S}_v , and temporal, \mathcal{T}_v , neighbourhoods for node v ($\mathcal{N}_v = \mathcal{S}_v \cup \mathcal{T}_v$). Each node, v , is associated with a latent state, h_v . Inference in this model consists of a message passing phase, and a final readout phase. In the message passing phase, messages for each node, m_v , are first computed by applying spatial and temporal message passing functions, M_s and M_t respectively, to all nodes in its neighbourhood as described in Eq. (1). An update function, U , then aggregates the received messages to update the latent state, h_v ,

$$m_v^{i+1} = \sum_{w \in \mathcal{S}_v} M_s(h_v^i, h_w^i; \theta_s^i) + \sum_{u \in \mathcal{T}_v} M_t(h_v^i, h_u^i; \theta_t^i) \quad (1)$$

$$h_v^{i+1} = U(m_v^{i+1}, h_v^i), \quad (2)$$

where θ denotes learnable function parameters. Intuitively, the state of a node, h_v , is updated by aggregating the messages passed from its neighbours. Finally, after $I \geq 1$ iterations of message passing, a readout function, R , uses the updated node features for the classification task of interest

$$y = R(\{h_v^i\}_{v \in \mathcal{G}}). \quad (3)$$

As illustrated in Fig. 2, our graph consists of a set of “Foreground” nodes, $\mathcal{F} = \{f_1, f_2, \dots, f_n\}$ and “Context” nodes, $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$ where n and m vary for each video. The “Foreground” nodes, \mathcal{F} , have readout and loss functions associated with them, and correspond to the objects that will subsequently be classified. For the task of action recognition, \mathcal{F} corresponds to bounding boxes of each actor in the keyframe. Whilst for scene graph prediction, \mathcal{F} refers to bounding boxes for all potential objects of interest. The “Context” nodes, \mathcal{C} , capture additional information extracted from the scene for relational reasoning. The representations of these nodes, and their spatial connectivity, are detailed next.

3.2. Spatial Model

The spatial connections in our graph model relationships between actors, objects and scene context in the same frame. For example, recognising an action such as “give object to a person” in the AVA dataset [17] requires understanding both people involved in the action and also the object being transferred.

To model such interactions, we first represent the foreground nodes of our model, \mathcal{F} , by extracting convolutional features from the last layer of the network, $\mathbf{X} \in \mathbb{R}^{t \times h \times w \times c}$, and using ROI-Align [19] followed by spatio-temporal pooling and a linear projection to obtain $f_i \in \mathbb{R}^d$.

We model scene context by considering the features from each spatial position in the feature map, \mathbf{X} , as a contextual node, c_i , in our graph. Note that these features are projected to \mathbb{R}^d . A similar representation was used by [53], which we refer to as an *implicit* object model as it enables the network to encode information about the scene and relevant objects without any extra supervision. This approach is also known as “grid features” in visual question answering [4, 24].

It is also possible to augment our set of contextual nodes with an *explicit* object representation by computing class-agnostic object proposals with a Region Proposal Network (RPN) [45]. We use an RPN pretrained on the OpenImages dataset [31] and obtain a d -dimensional feature from each proposal by using ROI-Align and a linear projection as for the foreground nodes. A similar idea of using explicit object representations in video has also been employed by [3] and [59], though not for spatio-temporal action recognition.

Messages are then computed and passed to foreground nodes from both foreground and contextual nodes in the graph. We only update the foreground nodes during message passing, as these are the nodes which are subsequently classified in the final Readout phase. Concretely, the spatial neighbourhood for each foreground node in our graph is

$$\mathcal{S}_v = \mathcal{F} \cup \mathcal{C} \quad \forall v \in \mathcal{F}. \quad (4)$$

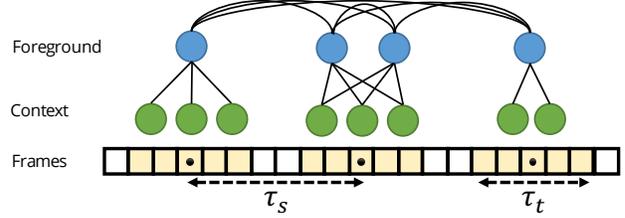


Figure 3: Illustration of temporal connectivity. Foreground (blue) and context (green) nodes have initial states computed from spatio-temporal features over τ_t frames (shaded in yellow), centred on a keyframe (denoted by a black circle). τ_s controls the distance between selected keyframes; τ_c is the total number of keyframes considered. Here, $\tau_c = 3$, $\tau_s = 7$ frames and $\tau_t = 5$ frames.

3.3. Temporal Model

We also include temporal connections in our graph to model long-range interactions between actors and objects. As shown in Fig. 1, understanding actions often requires reasoning about actors who are no longer visible in the current frame, thus requiring large temporal contexts.

We model these temporal interactions by connecting foreground nodes in keyframe t with all other foreground nodes in neighbouring keyframes $t' \in \mathcal{T}$. Concretely, we define the temporal context, τ_c , and temporal stride, τ_s , hyperparameters. As shown in Fig. 3, τ_c is the total number of number of keyframes in the video which we consider in our temporal graph. τ_s is the sampling rate at which we select keyframes, as $\tau_s \geq 1$ allows us to consider a wider temporal interval in a more computationally efficient manner. This is necessary to train the entire model end-to-end. Moreover, as each foreground feature node in the graph, f_i , is itself a spatio-temporal feature computed over a period of τ_t by a 3D CNN, selecting adjacent keyframes (and effectively setting $\tau_s = 1$ keyframe) could result in redundant information being captured by the temporal connections in the graph. Note that our definition of keyframe follows common datasets [17, 23] as a frame we aim to classify given surrounding temporal context.

More formally, we can describe the neighbourhood of each foreground node, v , for temporal message passing as

$$\mathcal{T}_v = \bigcup_{[t=-\tau_c/2]^{[\tau_c/2]}} \mathcal{F}^{t \cdot \tau_s}. \quad (5)$$

Here, we use the superscript to denote the temporal index, and without loss of generality consider $t = 0$ to be the centre keyframe, meaning that negative time indices correspond to frames in the past. We set τ_c to be an odd, positive integer to employ an equal-sized temporal window on either side of the centre keyframe.

We first perform spatial message passing, before passing messages temporally. This allows information from the context nodes in frame t' , $c_j^{t'}$, to efficiently propagate to a

foreground node f_i^t (where $t \neq t'$) via $f_j^{t'}$, as the foreground nodes are fully-connected temporally.

The following section now describes the messages that are passed along the graph described above.

3.4. Message Passing Functions

We first observe that the Non-local operator or self-attention [55, 58] can be considered as a message passing function in a fully-connected graph, where each node, h_v , is an element in the input feature map, \mathbf{H} , and the neighbourhood comprises of all other feature elements. And since self-attention is employed with a residual connection [58, 55] and layer normalisation [2, 55], the update function of Non-local when viewed as an MPNN is

$$U = \text{LN}(h_v + \text{Self-Attention}(\mathbf{H})). \quad (6)$$

Similar analysis [5, 13] has shown that Graph Attention Networks (GAT) [56], Relational Networks [46] and many other graph neural networks [47, 29, 35, 4] can also be interpreted as MPNNs. In this paper, we consider Non-local [58] and Graph Attention (GAT) [56] as message passing functions within the update function, Eq. (6).

Non-local We modify Non-local [58] to pass messages from all Foreground and Context nodes to only Foreground nodes (Eq. 4), as these are the nodes which are used for the final classification in the Readout phase,

$$\mathbf{M} = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\right)\mathbf{V} \quad (7)$$

$$\mathbf{Q} = \mathbf{A}\mathbf{W}_q \quad \mathbf{K} = [\mathbf{A}||\mathbf{C}]\mathbf{W}_k \quad \mathbf{V} = [\mathbf{A}||\mathbf{C}]\mathbf{W}_v.$$

Here, $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{C} \in \mathbb{R}^{m \times d}$ are matrices where each row is a foreground and context feature node respectively, \mathbf{M} is the $\mathbb{R}^{n \times d}$ matrix of messages received by each node in \mathbf{A} , $[\mathbf{A}||\mathbf{C}]$ denotes the concatenation of these matrices and \mathbf{W}_q , \mathbf{W}_k and \mathbf{W}_v are learnable $d \times d$ projection matrices.

Graph attention (GAT) The graph attention [56] message for a node v is computed as

$$m_v = \sigma\left(\sum_{j \in \mathcal{N}_v} \alpha_{ij} \mathbf{W}_a h_j\right) \quad (8)$$

$$\alpha_{ij} = \text{Softmax}\left(\sigma(w_b^\top [h_i || h_j])\right) \quad (9)$$

where σ is a ReLU non-linearity and \mathbf{W}_a and w_b are a learnable matrix and vector respectively.

Parallel messages It is also possible to compute multiple incoming messages in parallel for a node, h_v . When using multiple Non-local or Graph Attention functions, this corresponds to multi-headed attention [55, 56]. By viewing Non-local and GAT as message passing functions, we can also aggregate messages from a combination of these two. In these cases, we aggregate the messages using an attention-weighted convex combination, as performed in Eq. (9).

3.5. Readout Function

After $I \geq 1$ iterations of message passing, a readout function is applied on the Foreground nodes to obtain the final predictions. For action detection, the readout function is a linear classifier operating on each element of \mathcal{F} , where each foreground node, f_i , corresponds to the features of an actor in the keyframe. For scene graph prediction, the readout function consists of two classifiers: The first linear classifier predicts the object-class label of each foreground node. The second is a function of each pair of foreground nodes, and predicts the relationship label between them.

3.6. Discussion

We note that many previous structured models for video understanding can be considered as special cases of our proposed MPNN framework:

Girdhar *et al.* [14] only consider a spatial model, *i.e.* $\mathcal{N} = \mathcal{S}$, using implicit objects and Non-local [59] for message passing. ACRN [53] has the same graph structure, but uses Relational Networks [46] for message passing instead. LFB [60], in contrast, considers only a temporal model, *i.e.* $\mathcal{N} = \mathcal{T}$, using Non-local as the message passing function in a graph that is fully-connected in time. However, [60] do not consider a spatial model to capture interactions between actors in the keyframe. Zhang *et al.* [69], on the other hand, model both temporal and spatial connections. However, they effectively model three separate graphs: the first models actors in short, 3-second tubelets using GCN [29] for message passing. The other two graphs model actor-actor and actor-object relations using a message passing method similar to GAT [56], but using a hand-defined weighting function rather than a learned one as in GAT (Eq. 9). Wang *et al.* [59] also model a spatio-temporal graph, using GCN for message passing, and an explicit object representation in the spatial model. However, their temporal connections are only among adjacent frames, which does not allow information propagation between all frames in long sequences, in contrast to our model which is fully-connected temporally. With our unified framework, we study the effect of graph modelling design choices and show how we outperform previous work next in Sec. 4.

4. Experiments

We evaluate our method’s ability to model spatio-temporal interactions with experiments on spatio-temporal action detection and video scene graph classification.

Spatio-temporal action recognition We evaluate on AVA [17], the largest dataset for this task consisting of 15-minute video clips obtained from movies, and UCF101-24 [52], the previous standard benchmark for this task. AVA is labelled with atomic actions, where one person is typically performing multiple actions simultaneously, whilst ac-

Table 1: Experiments on Action Genome using a 3D ResNet 50 backbone. We ablate (a) different spatial message passing functions, (b) temporal connections in the graph and (c) iterations of message passing. (d) Comparison to existing methods which were reported by [23].

(a) Message passing functions		(b) Temporal graph structure		(c) Message passing iterations		(d) Comparison to existing methods				
	SGCls	Temporal parameters		Iterations	SGCls	SGCls		PredCls		
	R@20	τ_c	τ_s		R@20	R@20	R@50	R@10	R@20	
Baseline	48.9	Spatial only		1	51.1	MSDN [34]	44.0	47.2	-	-
Non-local in backbone	49.1	3	2	2	51.6	IMP [62]	44.1	47.4	-	-
		3	5	3	51.6	RelDN [67]	46.7	49.4	-	-
Non-local	50.4	3	7	5	51.8					
GAT	51.1	5	2			SlowFast (ResNet 50)	48.9	51.3	78.7	93.8
GAT + Non-local	51.3	5	5			Ours (ResNet 50)	53.8	56.0	79.3	94.2
		5	7							

tors perform only a single high-level action in UCF101-24. We follow standard protocol and evaluate using the Frame AP at an IoU threshold of 0.5 on both datasets. For AVA, we use v2.2 annotations for ablations, and either v2.1 or v2.2 annotations for fair comparisons with prior work.

Video scene graph prediction We evaluate on the recent Action Genome dataset [23], which adds scene graph annotations to Charades [48], on two scene graph tasks: scene graph classification (SGCls) and predicate classification (PredCls). Both tasks are evaluated using the standard Recall@K metric [38, 23] (R@K) which measures the fraction of ground truth relationship triplets (subject-predicate-object) that appear in the top K scoring predicted triplets. In SGCls, ground truth bounding box co-ordinates are given, and the aim is to predict their object classes, as well as relationship labels between pairs of objects. PredCls is simpler, as both bounding box co-ordinates and object classes are given, and only the relationship label must be predicted.

Implementation details We use the public implementation of SlowFast [11] as our baseline with the 3D ResNet 50 or ResNet 101 backbones [20], as it is the current state-of-the-art. The network is similar to Fast-RCNN [15] as it uses external region proposals to extract features from the last feature map of res_5 using ROI-Align [19]. These features are then spatio-temporally pooled and classified.

In our graph model, we initialise the internal states, h_i , of our Foreground nodes using these ROI-Aligned res_5 features. For our action detection experiments, the Foreground nodes in our graph correspond to bounding boxes of actors, and we use the same person detector as [11, 60] for our actor region proposals. For scene graph experiments, we use ground truth boxes of people and objects as our Foreground nodes as we evaluate scene graph classification (SGCls) and predicate classification (PredCls).

Unless otherwise specified, we use SlowFast 8×8 which corresponds to 32 input frames where the video is subsampled by a factor of 2. This means that each feature node, h_i , in the graph aggregates $\tau_t = 2.1$ seconds of temporal information on AVA as the videos are sampled at 30 frames per second (fps). As keyframes in AVA are defined at 1 fps, it means that we need to set $\tau_s \geq 2$ keyframes for the tem-

poral information captured by temporally-adjacent nodes in the graph to not overlap. For Action Genome, the Charades videos are sampled at 24 fps and the keyframes are on average 0.85 seconds, or 20.5 frames apart.

We train our network for 20 epochs using synchronous SGD on 8 GPUs and a total batch size of 64, initialising from a Kinetics-400 [28] pretrained model. The baseline model in all our experiments is SlowFast without any graph module trained in an identical manner. When training spatio-temporal graph models, we finetune for 10 epochs from a model trained with only a spatial graph and reduce the batch size by a factor of τ_c . We also “freeze” batch normalisation statistics as they have high variance for small batches [19]. Full training details are in the supplementary.

4.1. Experiments on video scene graph prediction

Spatial message passing We first ablate only the spatial component (Sec. 3.2) of our model in Tab. 1a. Graph Attention (GAT) [56] performs slightly better than Non-local [55, 58] as a message passing function, and we obtain a further small improvement by combining the two methods in parallel (Sec. 3.4). Overall, we improve upon our baseline model, which is SlowFast based on ResNet 50 without any graph modelling, by 2.4 points for R@20 on SGCls.

Another baseline is to insert a Non-Local layer [58] after the final res_5 layer and before the ROI-Align layer, since Non-local can also be viewed as a graph network (Sec. 3.4). In this case, the Foreground and Context nodes are not explicitly modelled as in our method and the overall performance is 2.2 points less than our method for the R@20 of SGCls. This suggests that explicit modelling of the Foreground nodes which are subsequently classified is important to performance. We note, however, that Non-Local [58] is typically employed earlier in a network (*i.e.* in res_3 or res_4 [58, 60, 11]) and can thus be seen as a complementary method to improve features learned by the network.

Temporal message passing Table 1b adds temporal connections (Sec. 3.3) to our model, using GAT for message passing, as it outperformed Non-local in Tab. 1a. We observe consistent improvements for a wide range of temporal contexts, τ_c , and strides, τ_s , showing the utility of modelling

Table 2: Analysis of message passing neighbourhood on AVA using SlowFast with a ResNet-50 backbone as the baseline. By varying the neighbourhood, we study the effect of object representations (implicit, explicit or none) and temporal connections. We report the Frame mAP on the three types of action classes in AVA.

Message passing neighbourhood	Pose	Human-Human	Human-Object	All
SlowFast baseline (none)	43.1	25.2	17.4	24.8
Actors only	43.2	27.0	17.8	25.6
Implicit objects only	43.4	26.7	18.0	25.7
Explicit objects only	43.0	26.7	17.8	25.5
Actor + Implicit	43.4	26.8	18.3	25.9
Actor + Implicit + Explicit	43.7	27.0	18.4	26.1
Spatio-temporal	43.8	27.5	19.9	27.0

temporal dynamics. We improve over a graph with only spatial connections by as much as 2.7 points for R@20.

For a temporal model, another baseline to consider is to simply increase the temporal information available to the spatial-only graph model by increasing the number of frames at the input. When tripling the number of input frames, the R@20 is 52.4, less than all variants of our temporal model with $\tau_c = 3$. Similarly, increasing input frames by a factor of 5, the R@20 is 52.8, less than all our temporal models with $\tau_c = 5$. This improvement demonstrates the benefits of message passing on an explicit temporal graph.

Iterations of message passing Table 1c shows that the model’s performance plateaus after 2 iterations of message passing, where parameters are not tied across iterations.

Comparison to state-of-the-art Finally, we compare to prior works in Tab. 1d which are single-image models evaluated by [23]. Our SlowFast, ResNet 50 3D baseline outperforms these, showing the importance of using spatio-temporal features for this task. As Action Genome contains videos of humans acting to scripts, there is temporal structure in the interactions between actors and objects in the scene (Fig. 6). Our final spatio-temporal graph structured model improves substantially upon this baseline by 4.9 and 4.7 points for the R@20 and R@50 for SGCLs respectively. Our improvements over the baseline for PredCIs are less, as this task is easier and the performance is saturated.

4.2. Experiments on action recognition

Graph structure and object representation Table 2 compares the effect of changing the neighbourhood that passes messages to each foreground node that is subsequently classified. We report the performance across the three types of action categories in AVA [17] – Pose, Human-Human and Human-Object – to show the effect that different object representations (implicit, explicit or neither as described in Sec. 3.2) and also temporal connections have. Our SlowFast baseline does not model any explicit graph and hence effectively performs no message passing. For

Table 3: Comparison on AVA. We report the Mean AP using v2.1 and v2.2 labels. All methods pretrained on Kinetics 400. “Multiscale” refers to averaging results over three scales [60].

Method	v2.1	v2.2
ACRN (S3D) [53]	17.4	–
Zhang <i>et al.</i> (R50) [69]	22.2	–
SlowFast baseline (R50)	24.5	24.8
Girdhar <i>et al.</i> (I3D) [14]	25.0	–
LFB (R50) [60]	25.8	–
Ours (R50)	26.5	27.0
Ours Multiscale (R50)	27.3	27.7
SlowFast baseline (R101)	26.3	26.7
LFB (R101) [60]	26.8	–
Ours (R101)	28.3	28.8
LFB Multiscale (R101) [60]	27.7	–
Ours Multiscale (R101)	29.5	30.0

Table 4: Comparison to state-of-the-art on UCF101-24. We report the Frame AP at 0.5 using the corrected annotations of [50].

Method	Modality	Mean AP
ACT [26]	RGB + Flow	69.5
Song <i>et al.</i> [51]	RGB + Flow	72.1
STEP [65]	RGB + Flow	75.0
Gu <i>et al.</i> [17]	RGB + Flow	76.3
MOC [36]	RGB + Flow	78.0
SlowFast (R50)	RGB	76.6
SlowFast (R101)	RGB	77.4
Ours (R50)	RGB	78.6
Ours (R101)	RGB	79.3

our AVA experiments, Foreground nodes in the graph correspond to bounding boxes of the actors in the scene, using the same person detections as [60, 11]. Thus, in this section, “Foreground” and “actor” nodes are used interchangeably.

We observe from the second row of Tab. 2 that passing messages only between actor nodes (and thus not modelling object interactions) provides an overall improvement of 0.8 points. The largest gain, as expected, are for Human-Human action classes, as these are the interactions modelled by passing messages between actor nodes in the graph. When passing messages from implicit context nodes to the actors (third row), we observe an improvement in all types of action classes, but primarily in Human-Human and Human-Object classes. This is because implicit context nodes encompass the entire feature map, and thus capture information about the whole scene (as also shown in Fig. 4).

Modelling objects explicitly with external regional proposals (fourth row) improves on the same action types as the implicit object model, but performs marginally worse. This suggests that our RPN trained on OpenImages [31] is unable to detect the objects that are most discriminative of AVA actions. As object annotations are not provided in AVA, it is not possible to evaluate the recall of our external region proposals on relevant objects. It is not clear if we should expect more improvement from explicit objects, because although previous works have considered implicit [14, 53, 69] and explicit representations [3, 59], we are not aware of any that have compared the two. Note that these explicit context nodes also model Human-Human interactions as our RPN is trained on people in OpenImages [31]. We obtain further improvements by combining messages from actors, implicit and explicit context nodes, as shown by the next two rows.

Finally, we evaluate our spatio-temporal model, which passes messages from actors, implicit- and explicit-context nodes with $\tau_c = 3$ and $\tau_s = 3$, corresponding to an overall temporal window of 8.5 seconds. This model performs the best, improving upon the baseline by 2.2 points, or a relative improvement of 8.9%. In particular, the temporal connections help to improve on Human-Object action classes.



(a) talk to (0.6), sit (0.92), watch person (0.81) (b) listen to person (0.9), watch person (0.89), sit (0.97) (c) carry/hold (1.0), sit(0.88)

Figure 4: Visualisation of spatial message passing when using implicit objects as detailed in Sec. 4.3. Note in (a) and (b) that the network focuses on the selected actor (denoted by the green box on the left) when the action is “talk to”, and on the person speaking to him when the action is “listen to person”. Attention weights, α_{ij} are colour-coded according to the bar on the far right. Best viewed in colour.



talk to (0.85), watch a person (0.82), sit(0.51)

Figure 5: Visualisation of temporal graph as detailed in Sec. 4.3. We show the neighbours for the actor depicted by the red box in the centre keyframe, colour-coding the respective boxes with their attention weights (α_{ij}). Predicted action scores of the red actor are below, and selected keyframes are $\tau_s = 2.1$ seconds apart. Notice that the model focuses on the people 2 keyframes (4.2 seconds) on either side of the actor to recognise that she is “talking to” and “watching” them.

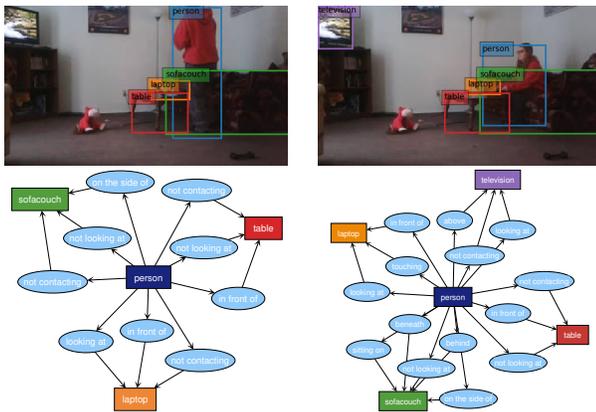


Figure 6: Example scene graph predictions for consecutive keyframes of Action Genome [23]. Classified bounding boxes are on top, and the corresponding scene graph (objects are shown as coloured rectangles, and relations with light-blue ovals) below.

AVA is a long-tailed dataset, and the most common failure modes of all variants of our model are the classes with few training examples. However, our model improves on both head and tail classes with respect to the baseline. This, and detailed per-class results, are in the supplementary.

State-of-the-art-comparison on AVA Table 3 compares to recent, published work on AVA. Our method builds upon SlowFast [11] as the base architecture, and our graph model shows substantial improvements with either a 3D ResNet 50 or ResNet 101 backbone. As discussed in Sec. 3.6, Girshar *et al.* [14] and LFB [60] can be considered as special cases of our graph model as they model only spatial and temporal edges respectively. Our method, which constructs a spatio-temporal graph, outperforms both. Note that we outperform LFB [60] across both ResNet 50 and ResNet 101 backbones. Zhang *et al.* [69] also model a spatio-temporal interactions, but employ three separate graphs with a hand-crafted aggregation function which we also out-

perform with our single, coherent spatio-temporal graph.

State-of-the-art comparison on UCF101-24 Table 4 shows that we outperform recent, published work on UCF101-24 [52], using either a 3D ResNet 50 or ResNet 101 backbone. We do so without using optical flow as an additional input modality, showing that our network can capture temporal information without it. Moreover, the improvements obtained from our graph model are consistent with our results on the Action Genome and AVA datasets.

4.3. Qualitative results

We visualise the spatial and temporal messages received by an actor node in the graph, on the AVA dataset, when using GAT as the message passing function. In particular, we visualise the attention weights, α_{ij} in (9), on each implicit object node (Fig. 4) and actor in neighbouring keyframes (Fig. 5), for a given actor. We observe that the network places more weight on graph nodes that are intuitively consistent with the final action prediction. For example, in Fig. 4, the model focuses on the actor’s face when his action is “talk to” and the other person’s face when his action is “listen to”. Finally, Fig. 6 shows example scene graph predictions. Further examples are in the supplementary.

5. Conclusion and Future Work

We have presented a novel spatio-temporal graph neural network framework to explicitly model interactions between actors, objects and their environment. Our formulation can model objects either implicitly or explicitly, and generalises previous structured models for video understanding [14, 53, 60, 69]. Using our versatile approach, we have achieved state-of-the-art results on two diverse tasks across three datasets. Future work remains to harness explicit object representations more effectively on AVA.

References

- [1] Anurag Arnab, Chen Sun, Arsha Nagrani, and Cordelia Schmid. Uncertainty-aware weakly supervised action detection from untrimmed videos. In *ECCV*, 2020. 2
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. In *arXiv preprint arXiv:1607.06450*, 2016. 5
- [3] Fabien Baradel, Natalia Neverova, Christian Wolf, Julien Mille, and Greg Mori. Object level visual reasoning in videos. In *ECCV*, 2018. 1, 2, 4, 7
- [4] Peter Battaglia, Razvan Pascanu, Matthew Lai, and Danilo Jimenez Rezende. Interaction networks for learning about objects, relations and physics. In *NeurIPS*, 2016. 3, 4, 5
- [5] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. In *arXiv preprint arXiv:1806.01261*, 2018. 2, 3, 5
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 1
- [8] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. A²-nets: Double attention networks. In *NeurIPS*, 2018. 2
- [9] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Yan Shuicheng, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *CVPR*, 2019. 2
- [10] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *CVPR*, 2017. 3
- [11] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 6, 7, 8
- [12] Christoph Feichtenhofer, Axel Pinz, and Richard Wildes. Spatiotemporal residual networks for video action recognition. In *NeurIPS*, 2016. 2
- [13] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *ICML*, 2017. 1, 2, 3, 5
- [14] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *CVPR*, 2019. 1, 2, 3, 5, 7, 8
- [15] Ross Girshick. Fast r-cnn. In *ICCV*, 2015. 6
- [16] Georgia Gkioxari, Ross Girshick, and Jitendra Malik. Contextual action recognition with r* cnn. In *ICCV*, 2015. 2
- [17] 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. 1, 2, 4, 5, 7
- [18] Abhinav Gupta, Aniruddha Kembhavi, and Larry S Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *PAMI*, 31(10), 2009. 2
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 4, 6
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 6
- [21] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 2
- [22] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *CVPR*, 2016. 2
- [23] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action genome: Actions as compositions of spatio-temporal scene graphs. In *CVPR*, 2020. 2, 3, 4, 6, 7, 8
- [24] Huaizu Jiang, Ishan Misra, Marcus Rohrbach, Erik Learned-Miller, and Xinlei Chen. In defense of grid features for visual question answering. In *CVPR*, 2020. 4
- [25] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *CVPR*, 2015. 3
- [26] Vicky Kalogeiton, Philippe Weinzaepfel, Vittorio Ferrari, and Cordelia Schmid. Action tubelet detector for spatio-temporal action localization. In *ICCV*, 2017. 2, 7
- [27] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 2
- [28] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. In *arXiv preprint arXiv:1705.06950*, 2017. 6
- [29] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 2, 3, 5
- [30] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123(1), 2017. 3
- [31] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020. 4, 7
- [32] Ivan Laptev. On space-time interest points. *IJCV*, 64(2-3), 2005. 2
- [33] Yikang Li, Wanli Ouyang, Xiaogang Wang, and Xiao'ou Tang. Vip-cnn: Visual phrase guided convolutional neural network. In *CVPR*, 2017. 3
- [34] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. 6

- [35] Yujia Li, Daniel Tarlow, Marc Brockschmidt, and Richard Zemel. Gated graph sequence neural networks. In *ICLR*, 2016. 5
- [36] Yixuan Li, Zixu Wang, Limin Wang, and Gangshan Wu. Actions as moving points. In *ECCV*, 2020. 7
- [37] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 2
- [38] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, 2016. 6
- [39] Chih-Yao Ma, Asim Kadav, Iain Melvin, Zsolt Kira, Ghassan AlRegib, and Hans Peter Graf. Attend and interact: Higher-order object interactions for video understanding. In *CVPR*, 2018. 2
- [40] Marcin Marszalek, Ivan Laptev, and Cordelia Schmid. Actions in Context. In *CVPR*, 2009. 2
- [41] Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, and Trevor Darrell. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *CVPR*, 2020. 2
- [42] Aude Oliva and Antonio Torralba. The role of context in object recognition. *Trends in cognitive sciences*, 11(12), 2007. 2
- [43] Siyuan Qi, Wenguan Wang, Baoxiong Jia, Jianbing Shen, and Song-Chun Zhu. Learning human-object interactions by graph parsing neural networks. In *ECCV*, 2018. 2
- [44] Andrew Rabinovich, Andrea Vedaldi, Carolina Galleguillos, Eric Wiewiora, and Serge Belongie. Objects in context. In *ICCV*, 2007. 2
- [45] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 1, 2, 3, 4
- [46] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *NeurIPS*, 2017. 5
- [47] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1), 2008. 2, 5
- [48] Gunnar A. Sigurdsson, G. Varol, X. Wang, Ali Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016. 1, 3, 6
- [49] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014. 2
- [50] Gurkirt Singh, Suman Saha, Michael Sapienza, Philip HS Torr, and Fabio Cuzzolin. Online real-time multiple spatiotemporal action localisation and prediction. In *ICCV*, 2017. 2, 7
- [51] Lin Song, Shiwei Zhang, Gang Yu, and Hongbin Sun. Tacnet: Transition-aware context network for spatio-temporal action detection. In *CVPR*, 2019. 7
- [52] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. In *arXiv preprint arXiv:1212.0402*, 2012. 2, 5, 8
- [53] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *ECCV*, 2018. 1, 2, 3, 4, 5, 7, 8
- [54] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 2
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2, 5, 6
- [56] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018. 3, 5, 6
- [57] Heng Wang, Alexander Kläser, Cordelia Schmid, and Cheng-Lin Liu. Dense trajectories and motion boundary descriptors for action recognition. *IJCV*, 103(1), 2013. 2
- [58] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 2, 5, 6
- [59] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018. 1, 2, 3, 4, 5, 7
- [60] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019. 1, 2, 3, 5, 6, 7, 8
- [61] Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. Pay less attention with lightweight and dynamic convolutions. In *ICLR*, 2019. 2
- [62] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, 2017. 3, 6
- [63] Huijuan Xu, Lizhi Yang, Stan Sclaroff, Kate Saenko, and Trevor Darrell. Spatio-temporal action detection with multi-object interaction. In *arXiv preprint arXiv:2004.00180*, 2020. 2
- [64] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *ECCV*, 2018. 3
- [65] Xitong Yang, Xiaodong Yang, Ming-Yu Liu, Fanyi Xiao, Larry S Davis, and Jan Kautz. Step: Spatio-temporal progressive learning for video action detection. In *CVPR*, 2019. 7
- [66] Bangpeng Yao and Li Fei-Fei. Modeling mutual context of object and human pose in human-object interaction activities. In *CVPR*, 2010. 2
- [67] Ji Zhang, Kevin J Shih, Ahmed Elgammal, Andrew Tao, and Bryan Catanzaro. Graphical contrastive losses for scene graph parsing. In *CVPR*, 2019. 6
- [68] Li Zhang, Dan Xu, Anurag Arnab, and Philip HS Torr. Dynamic graph message passing networks. In *CVPR*, 2020. 2
- [69] Yubo Zhang, Pavel Tokmakov, Martial Hebert, and Cordelia Schmid. A structured model for action detection. In *CVPR*, 2019. 1, 2, 3, 5, 7, 8
- [70] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018. 2