

# Object-Region Video Transformers

Roei Herzig<sup>1</sup> Elad Ben-Avraham<sup>1</sup> Karttikeya Mangalam<sup>2</sup> Amir Bar<sup>1</sup> Gal Chechik<sup>3</sup>  
Anna Rohrbach<sup>2</sup> Trevor Darrell<sup>2</sup> Amir Globerson<sup>1</sup>

<sup>1</sup>Tel Aviv University, <sup>2</sup>UC Berkeley, <sup>3</sup>Bar-Ilan University, NVIDIA Research




Figure 1. Our ORViT model incorporates object information into video transformer layers. The figure shows the standard (uniformly spaced) transformer patch-tokens in **blue**, and object-regions corresponding to detections in **orange**. In ORViT any temporal patch-token (e.g., the patch in **black** at time  $T$ ) attends to all patch tokens (**blue**) and region tokens (**orange**). This allows the new patch representation to be informed by the objects. Our method shows strong performance improvement on multiple video understanding tasks and datasets, demonstrating the value of a model that incorporates object representations into a transformer architecture.

## Abstract

Recently, video transformers have shown great success in video understanding, exceeding CNN performance; yet existing video transformer models do not explicitly model objects, although objects can be essential for recognizing actions. In this work, we present *Object-Region Video Transformers (ORViT)*, an object-centric approach that extends video transformer layers with a block that directly incorporates object representations. The key idea is to fuse object-centric representations starting from early layers and propagate them into the transformer-layers, thus affecting the spatio-temporal representations throughout the network. Our ORViT block consists of two object-level streams: appearance and dynamics. In the appearance stream, an “Object-Region Attention” module applies self-attention over the patches and object regions. In this way, visual object regions interact with uniform patch tokens and enrich them with contextualized object information. We further model object dynamics via a separate “Object-Dynamics Module”, which captures trajectory interactions, and show how to integrate the two streams. We evaluate our model on four tasks and five datasets: compositional and few-shot action recognition on *SomethingElse*, spatio-temporal action detection on AVA, and standard action recognition on *Something-Something V2*, *Diving48* and

*Epic-Kitchen100*. We show strong performance improvement across all tasks and datasets considered, demonstrating the value of a model that incorporates object representations into a transformer architecture. For code and pretrained models, visit the project page at <https://roeiherz.github.io/ORViT/>

## 1. Introduction

Consider the simple action of “Picking up a coffee cup” in Figure 1. Intuitively, a human recognizing this action would identify the hand, the coffee cup and the coaster, and perceive the upward movement of the cup. This highlights three important cues needed for recognizing actions: what/where are the objects? how do they interact? and how do they move? The above perception process allows easy generalization to different compositions of actions. For example, the process of “picking up a knife” shares some of the components with “picking up a coffee cup”, namely, the way the object and hand move together. More broadly, representing image semantics using objects facilitates compositional understanding, because many perceptual components remain similar when one object is swapped for another. Thus, a model that captures this compositional aspect potentially requires less examples to train.

It seems intuitively clear that machine vision models

should also benefit from such object-focused representations, and indeed this has been explored in the past [29, 69] and more recently by [3, 63, 83], who utilize bounding boxes for various video understanding tasks. However, the central question of what is the best way to process objects information still remains. Most object-centric methods to video understanding take a post-processing approach. Namely, they compute object descriptors using a backbone and then re-estimate those based on other objects via message passing or graph networks without propagating the object information into the backbone. Unlike these approaches, we argue that objects should influence the spatio-temporal representations of the scene throughout the network, starting from the early layers (i.e., closer to the input). We claim that self-attention in video transformers is a natural architecture to achieve this result by enabling the attention to incorporate objects *as well as* salient image regions.

Video transformers have recently been introduced as powerful video understanding models [2, 7, 30, 65], motivated by the success of transformers in language [17] and vision [10, 19]. In these models, each video frame is divided into patches, and a self-attention architecture obtains a contextualized representation for the patches. However, this approach has no explicit representation of objects. Our key observation is that self-attention can be applied jointly to object representations and spatio-temporal representations, thus offering an elegant and straightforward mechanism to enhance the spatio-temporal representations by the objects.

Motivated by the above, our key goal in this paper is to explicitly fuse object-centric representations into the spatio-temporal representations of video-transformer architectures [2], and do so throughout the model layers, starting from the earlier layers. We propose a general approach for achieving this by adapting the self-attention block [19] to incorporate object information. The challenge in building such an architecture is that it should have components for modeling the appearance of objects as they move, the interaction between objects, and the dynamics of the objects (irrespective of their visual appearance). An additional desideratum is that video content outside the objects should not be discarded, as it contains important contextual information. In what follows, we show that a self-attention architecture can be extended to address these aspects. Our key idea is that object regions can be introduced into transformers in a similar way to that of the regular patches, and dynamics can also be integrated into this framework in a natural way. We refer to our model as an “Object-Region Video Transformer” (or ORViT).

We introduce a new ORViT block, which takes as input bounding boxes and patch tokens (also referred to as spatio-temporal representations) and outputs refined patch tokens based on object information. Within the block, the information is processed by two separate object-level streams:

an “Object-Region Attention” stream that models appearance and an “Object-Dynamics Module” stream that models trajectories.<sup>1</sup> The appearance stream first extracts descriptors for each object based on the object coordinates and the patch tokens. Next, we append the object descriptors to the patch tokens, and self-attention is applied to all these tokens jointly, thus incorporating object information into the patch tokens (see Figure 1). The trajectory stream only uses object coordinates to model the geometry of motion and performs self-attention over those. Finally, we re-integrate both streams into a set of refined patch tokens, which have the same dimensionality as the input to our ORViT block. This means that the ORViT block can be called repeatedly. See Figure 2 and Figure 4 for visualizations.

We evaluate ORViT on several challenging video-understanding tasks: compositional and few-shot action recognition on SomethingElse [63], where bounding boxes are given as part of the input; spatio-temporal action detection on AVA [28], where the boxes are obtained via an off-the-shelf detector that was provided by previous methods; and in a standard action recognition task on Something-Something V2 [26], Diving48 [54] and Epic-Kitchen100 [16], where we use class-agnostic boxes from an off-the-shelf detector. Through extensive empirical study, we show that integrating the ORViT block into the video transformer architecture leads to improved results on all tasks. These results confirm our hypothesis that incorporating object representations starting from early layers and propagating them into the spatio-temporal representations throughout the network, leads to better performance.

## 2. Related Work

**Object-centric models.** Recently object-centric models have been successfully applied in many computer vision applications: visual relational reasoning [5, 6, 35, 40, 49, 66, 89, 95], representation learning [91], video relation detection [55, 76], vision and language [13, 52, 53, 77], human-object interactions [23, 45, 87], and even image generation [33, 42]. The advances and the success of object-centric models in these domains inspired various video-based tasks, such as action localization [64], video synthesis [4], and action recognition [99]. The latter was the focus of varied recent works that designed different object interactions approaches for convolutional models. A line of works [25, 70, 75] focused on capturing spatial object interactions while ignoring the temporal interactions. STRG [83] and ORN [5] used spatio-temporal interactions with two consecutive frame interactions, while STAG [34] considered long-range temporal interaction. Last, Unified [3] tried to generalize all these models and propose long

<sup>1</sup>Our focus is different from papers on two-stream models in vision, that are not object-centric (see Sec. 2).

spatio-temporal object interactions. While all these works focused solely on interactions of visual appearance information, recently STIN [63] introduced an object-centric model based on object trajectories by modeling bounding box movement. Our ORViT approach directly combines object appearance, object trajectories, and the overall video, by mapping all computations to spatio-temporal patch tokens. This is particularly natural to do in a transformer framework, as we show here, and results in state-of-the-art performance.

**Transformers in action recognition.** Ranging from the early works that employ optical flow based features [20], to recent transformer based approaches [30], a wide variety of approaches have been proposed for action recognition. In broad brushstrokes, the proposed approaches have evolved from using temporal pooling for extracting features [44] to using recurrent networks [18, 94], through to 3D spatio-temporal kernels [11, 41, 56, 78, 80–82], and two-stream networks that capture complementary signals (e.g., motion and spatial cues [21, 22, 73]). Unlike these approaches, our work uses two separate object-level streams to leverage object-centric information. In parallel to developments in video understanding, Vision Transformers [19, 79] propose a new approach to image recognition by discarding the convolutional inductive bias entirely and instead employing self-attention operations. Specialized video models such as TimeSformer [7], ViViT [2], Mformer (MF) [65] and MViT [30] form the latest epoch in action recognition models. By generalizing the vision transformers to the temporal domain through the use of spatio-temporal attention, the obtained video transformers are very competitive with their convolutional counterparts both in terms of performance as well scaling behaviour to large data. However, none of the video transformer models leverage object cues, a persistent shortcoming that we aim to address in ORViT.

We also note that [85] adopts a similar object-centric approach to video understanding. However, our work differs conceptually since [85] models only the object parts, which is similar to the STRG baselines we consider in the paper. On the other hand, our work introduces objects into transformer layers while keeping the entire spatio-temporal representation. Also, [85] pretrains the transformer in a self-supervised fashion on a large dataset (MovieClips), and therefore its empirical results cannot be directly compared to models that are not pretrained in this manner.

**Spatio-temporal action detection.** The task of action detection requires temporally localizing the action start and end times. A wide variety of methods have been proposed for it, such as actions modeling [1, 61, 96], temporal convolutions [51, 72], boundaries modeling [57, 58], attention [71, 92], structure utilization [93, 98], detection based methods [12, 88], end-to-end approaches [9, 15, 24, 39], recurrent neural networks [62, 74, 90], and even using lan-

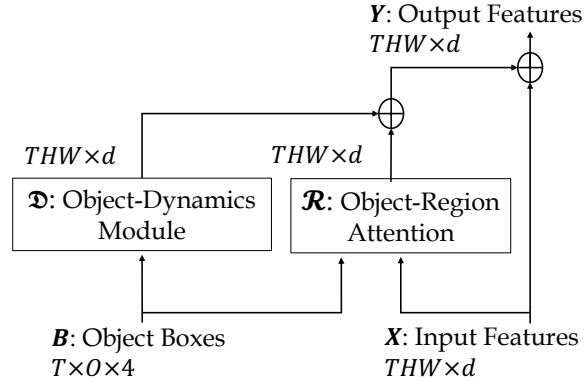


Figure 2. An ORViT block. The input patch-tokens  $X$  and boxes  $B$  are used as input to the “Object-Region Attention” and “Object-Dynamics Module” components. Each component outputs a  $THW \times d$  tensor and the two tensors are summed to obtain new patch tokens  $Y$ .

guage [68, 100]. Recently, the new MViT [30] model showed promising results on action localization in the AVA dataset [27]. However, it does not explicitly model objects, and we show that an ORViT version of MViT indeed improves performance.

### 3. The ORViT model

We next present the Object-Region Video Transformer (ORViT) model, which explicitly models object appearance and trajectories within the transformer architecture. We begin by reviewing the video transformer architecture, which our model extends upon, in Section 3.1, and present ORViT in Section 3.2. A high-level overview of ORViT is shown in Figure 2 and detailed in Figure 4. Briefly, ORViT repeatedly refines the patch token representations by using information about both the appearance and movement of objects.

#### 3.1. The Video Transformer Architecture

Video transformers [2, 7, 30] extend the Vision Transformer model to the temporal domain. Similar to vision transformers, the input is first “patchified” but with temporally extended 3-D patches instead of 2-D image patches producing a down-sampled tensor  $X$  of size  $T \times H \times W \times d$ . Then, spatio-temporal position embeddings are added for providing location information. Finally, a classification token (CLS) is appended to  $X$ , resulting in  $THW + 1$  tokens in  $\mathbb{R}^d$ , to which self-attention are applied repeatedly to produce the final contextualized CLS feature vector.<sup>2</sup>

#### 3.2. The ORViT Block

There are two inputs to the ORViT block. The first is the output of the preceding transformer block, represented as a set of spatio-temporal tokens  $X \in \mathbb{R}^{THW \times d}$ . The second

<sup>2</sup>In what follows, we omit the count of the CLS feature for brevity.

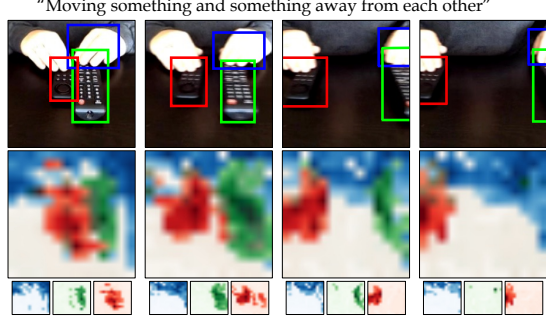


Figure 3. We visualize the attention allocated to the object tokens in the ORViT block (red, green, and blue) in each frame of a video describing “moving two objects away from each other”. It can be seen that each of the two “remote control” objects affects the patch-tokens in its region, whereas the hand has a broader map. For more visualizations, please see Section E in supplementary.

input is a set of bounding boxes for objects across time<sup>3</sup>, denoted by  $B \in \mathbb{R}^{TO \times 4}$ . The output of the ORViT block is a set of refined tokens  $Y \in \mathbb{R}^{THW \times d}$  contextualized with object-centric information. Thus, the ORViT block can be viewed as a token representation refining mechanism using the object-level information.

As mentioned, we argue that the key cues for recognizing actions in videos are: the objects in the scene, their interactions, and their movement. To capture these cues, we design the ORViT block with the following two object-level streams. The first stream models the appearance of objects, and their interactions. We refer to it as “Object-Region Attention” and denote it by  $\mathcal{R}$ . The second “Object-Dynamics Module” stream (denoted by  $\mathcal{D}$ ) models the interactions between trajectories, independently of their appearance. Importantly, the output of each of the streams is  $THW$  token vectors, which can also be interpreted as refined patch representations based on each source of information.

The  $\mathcal{D}$  stream only models object dynamics, and thus only uses bounding boxes  $B$  as input. We therefore denote its output by  $\mathcal{D}(B)$ . The stream  $\mathcal{R}$  models appearance and thus depends on both the token representation  $X$ , and the bounding boxes  $B$ , and produces  $\mathcal{R}(X, B)$ . The final output of the ORViT block  $Y$  is simply formed by the sum of the two streams and an input residual connection:

$$\begin{aligned} Y' &:= \mathcal{R}(X, B) + \mathcal{D}(B) + X \\ Y &:= Y' + \text{MLP}(\text{LN}(X)) \end{aligned} \quad (1)$$

where  $\text{LN}$  denotes a LayerNorm operation. Next, we elaborate on the two components separately.

**Object-Region Attention.** The goal of this module is to extract information about each object and use it to refine the patch tokens. This is done by using the object regions

<sup>3</sup> $O$  represents the maximum number of objects in the training set. If a clip has less than  $O$  boxes, we pad the remaining embeddings with zeros.

to extract descriptor vectors per region from the input tokens, resulting in  $TO$  vectors in  $\mathbb{R}^d$ , which we refer to as object tokens. These vectors are then concatenated with the  $THW$  patch tokens and serve as the keys and values, while the queries are only the patch tokens. Finally, the output of the block is  $THW$  patch tokens. Thus, the key idea is to fuse object-centric information into spatio-temporal representations. Namely, inject the  $TO$  object region tokens into  $THW$  patch tokens. An overview of our approach is depicted in Figure 4. We provide further details below.

Given the patch token features  $X$  and the boxes  $B$ , our first goal is to obtain vector descriptors in  $\mathbb{R}^d$  per object and frame. The natural way to do this is via an RoIAlign [31] layer, which uses the patch tokens  $X$  and box coordinates  $B$  to obtain object region crops. This is followed by max-pooling and an MLP to obtain the final object representation in  $\mathbb{R}^d$ :

$$\mathcal{O} := \text{MLP}(\text{MaxPool}(\text{RoIAlign}(X, B))) \quad (2)$$

Since this is done per object and per frame, the result is  $OT$  vectors in  $\mathbb{R}^d$  (i.e.,  $\mathcal{O} \in \mathbb{R}^{TO \times d}$ ). Importantly, this extraction procedure is performed in *each instance of an ORViT block*, so that it will produce different object tokens at each layer. We also add positional embeddings but leave the details to Section B.1 in supplementary.

At this point, we would like to allow the object tokens to refine the patch tokens. We concatenate the object tokens  $\mathcal{O}$  with the patch tokens  $X$ , resulting in  $C \in \mathbb{R}^{T(HW+O) \times d}$ . Next  $C$  and  $X$  are used to obtain queries, keys and values as follows:

$$\begin{aligned} Q &:= XW_q \quad K := CW_k \quad V := CW_v \\ \text{Where } W_q, W_k, W_v &\in \mathbb{R}^{d \times d} \end{aligned} \quad (3)$$

Finally, there are several ways to perform spatio-temporal self-attention (e.g., joint and divided attention over space and time, or the recently introduced trajectory attention [65]). We use trajectory attention because it performs well empirically. We compare different self-attention versions in Table 5c in supplementary. Figure 3 also visualizes the “Object-Region Attention” learned by our model.

**Object-Dynamics Module.** To model object dynamics, we introduce a component that only considers the boxes  $B$ . We first encode each box via its center coordinate, height and width, and apply an MLP to this vector to obtain a vector in  $\mathbb{R}^d$ . Applying this to all boxes results in  $\tilde{L} \in \mathbb{R}^{TO \times d}$ . Next we add a learnable object-time position embedding  $\tilde{P} \in \mathbb{R}^{TO \times d}$ , resulting in  $\tilde{B} := \tilde{L} + \tilde{P}$ . We refer to this as the “Coordinate Embedding” step in Figure 4. Its output can be viewed as  $TO$  tokens in  $\mathbb{R}^d$ , and we apply self-attention to those as follows:  $\text{Attention}_{\mathcal{D}}(\tilde{Q}, \tilde{K}, \tilde{V}) := \text{Softmax}\left(\frac{\tilde{Q}\tilde{K}^T}{\sqrt{d_k}}\right)\tilde{V}$ , where:  $\tilde{Q} := \tilde{B}W_{\tilde{q}}$ ,  $\tilde{K} := \tilde{B}W_{\tilde{k}}$ ,  $\tilde{V} := \tilde{B}W_{\tilde{v}}$  and  $W_{\tilde{q}}, W_{\tilde{k}}, W_{\tilde{v}} \in \mathbb{R}^{d \times d}$ . The self-attention

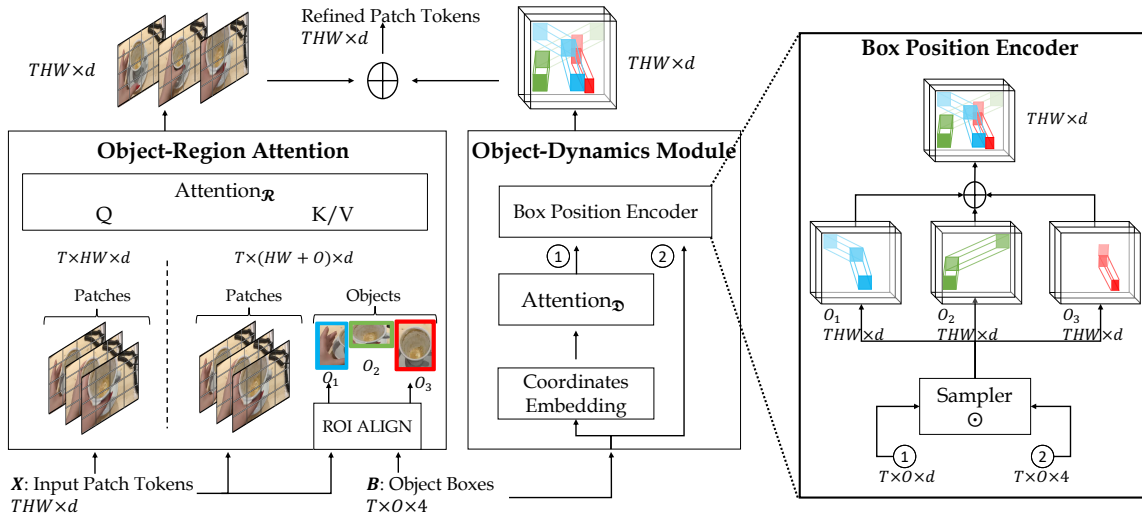


Figure 4. **ORViT Block architecture.** The block consists of two object-level streams: an “Object-Region Attention” that models appearance, and an “Object-Dynamics Module” that models trajectories. The two are combined to produce new patch tokens. The “Box Position Encoder” maps the output of the trajectory stream to the dimensional of patch tokens.

output is in  $\mathbb{R}^{TO \times d}$ . Next, we would like to transform the objects with a  $T \times d$  vector into a spatial volume of  $THW \times d$ . This is done using the Box Position Encoder described below.

**Box Position Encoder.** The returned features of the ORViT model should have the same dimensions as the input, namely  $THW \times d$ . Thus, our main challenge is to project the object embeddings into spatial dimensions, namely  $TO \times d$  into  $THW \times d$ . The naive approach would be to ignore the boxes by expanding every object with vector  $T \times d$  into  $THW \times d$ . However, since the object trajectories contain their space-time location, a potentially better way to do it would consider the object locations. Hence, for each object with corresponding  $T \times d$  tokens, we generate a spatial feature  $HW \times d$  by placing the object representation vector according to the matching bounding box coordinates using a bilinear interpolation sampler operation [38, 42].<sup>4</sup> Finally, the output in  $HW \times d$  is the sum of all objects from all frames representing the coarse trajectory of the object in spatial dimensions. The process is shown in Figure 4 (right). We show empirically that this approach is better than the naive approach described above.

### 3.3. The ORViT model

We conclude by explaining how to integrate ORViT into transformer-based video models. The advantage of ORViT is that it takes as input the standard spatio-temporal tokens in  $\mathbb{R}^{THW \times d}$  and outputs a refined version of those with the same dimensions. Thus, it acts as a standard transformer layer in terms of input and output, and one can take any transformer and simply add ORViT layers to it. This is important since it highlights that ORViT can eas-

<sup>4</sup>Features outside of an object region are set to zeros.

ily leverage any video transformer pretrained model, obviating the need for pretraining ORViT. We experiment with three video transformer models: TimeSformer [7], Mformer (MF) [65], and MViT [30]. We show that for these, using ORViT layers improves performance. The only design choice is which layers to apply ORViT to, while the training methodology remains. We found that it is very important to apply it in early layers while repeated applications keep propagating the information throughout the network. Since the RoIAlign extracts object representations from spatio-temporal representations in each ORViT layer, multiple ORViT layers allow the model to consider different object representations throughout the network. In our experiments we apply it in layers 2, 7, 11, replacing the original layers without adding depth to the baseline model.

## 4. Experiments

We evaluate our ORViT block on several video understanding benchmarks. Specifically, we consider the following tasks: Compositional Action Recognition (Section 4.1), Spatio-Temporal Action Detection (Section 4.2) and Action Recognition (Section 4.3).

**Datasets.** We experiment with the following datasets: (1) **Something-Something v2 (SSv2)** [26] is a dataset containing 174 action categories of common human-object interactions. (2) **SomethingElse** [63] exploits the compositional structure of SSv2, where a combination of a verb and a noun defines an action. We follow the official compositional split from [63], which assumes the set of noun-verb pairs available for training is disjoint from the set given at test time. (3) **Atomic Visual Actions (AVA)** [28] is a benchmark for human action detection. We report Mean Average Precision (mAP) on AVA-V2.2. (4) **Epic Kitchens 100 (EK100)** [16]

Model	Boxes	Compositional		Base		Few-Shot	
		Top-1	Top-5	Top-1	Top-5	5-Shot	10-Shot
I3D [11]	✗	42.8	71.3	73.6	92.2	21.8	26.7
SlowFast [21]	✗	45.2	73.4	76.1	93.4	22.4	29.2
TimeSformer [7]	✗	44.2	76.8	79.5	95.6	24.6	33.8
MF [65]	✗	60.2	85.8	82.8	96.2	28.9	33.8
STRG (\w SF) [83]	✓	52.3	78.3	75.4	92.7	24.8	29.9
STIN (\w SF) [63]	✓	54.6	79.4	77.4	95.0	23.0	33.4
MF+STRG+STIN	✓	62.3	86.0	83.7	96.8	29.8	36.5
<b>ORViT MF(ours)</b>	✓	<b>69.7</b>	<b>91.0</b>	<b>87.1</b>	<b>97.6</b>	<b>33.3</b>	<b>40.2</b>

Table 1. **Compositional and Few-Shot Action Recognition** on the “SomethingElse” dataset.

contains 700 egocentric videos of kitchen activities. This dataset includes noun and verb classes, and we report verb, noun, and action accuracy, where the highest-scoring verb and noun pair constitutes an action label. (5) **Diving48** [54] contains 48 fine-grained categories of diving activities.

**Baselines.** In the experiments, we compare ORViT to several models reported in previous work for the corresponding datasets. These include non-transformer approaches (e.g., *I3D* [11] and *SlowFast* [21]) as well as state-of-the-art transformers (TimeSformer, Mformer (MF), and MViT). We also cite results for two object-centric models: *STIN* [63] which uses boxes information, and the Space-Time Region Graph (STRG) model [83] which extracts I3D features for objects and runs a graph neural network on those. Both STIN and STRG use the same input information as ORViT. Finally, we implement an object-centric transformer baseline combining STRG and STIN: we use the MF final patch tokens as input to the STRG model, resulting in STRG feature vector, and concatenate it to the STIN feature vector and the MF’s CLS token. We refer to this as *MF+STRG+STIN*.

**Implementation Details.** ORViT is implemented in PyTorch, and the code will be released in our project page. Our training recipes and code are based on the MViT, MF, and TimeSformer code published by the authors. For all tasks and datasets, we use SORT [8, 43] for multi-object tracking to find correspondence between the objects in different frames (no training data is required), see Section A.1 in supp. We set the number of objects to 4 in SSv2 and EK100, 6 in AVA, and 10 in Diving48. These numbers were chosen by taking the max number of objects per video (as induced by the tracker) across all videos in the training set.

#### 4.1. Compositional & Few-Shot Action Recognition

Several video datasets define an action via a combination of a verb (action) and noun (object). In such cases, it is more challenging to recognize combinations that were not seen during training. This “compositional” setting was explored in the “SomethingElse” dataset [63], where verb-noun combinations in the test data do not occur in the training data. The split contains 174 classes with 54,919/54,876 videos for training/validation. This setting is of particular

Model	Boxes	Pretrain	mAP
SlowFast [21], $4 \times 16$ , R50	✓	K400	21.9
SlowFast [21], $8 \times 8$ , R50	✓	K400	22.7
SlowFast [21], $8 \times 8$ , R101	✓	K400	23.8
MViT-B [30], $16 \times 4$	✓	K400	25.5
MViT-B [30], $32 \times 3$	✓	K400	27.3
<b>ORViT MViT-B, <math>16 \times 4</math> (Ours)</b>	✓	K400	<b>26.6 (+1.1)</b>
<b>ORViT MViT-B, <math>32 \times 3</math> (Ours)</b>	✓	K400	<b>28.0 (+0.7)</b>

Table 2. **Spatio-temporal Action Detection** on AVA-V2.2.

relevance for object-centric models like ORViT, which can potentially better handle compositional actions. The dataset contains annotated bounding boxes that are commonly used as additional input to the model; this allows us to perform a fair comparison against previous methods [63]. We also evaluate on the few-shot compositional action recognition task in [63] (see Section A.6 in supplementary for details).

Table 1 reports the results for these tasks. ORViT outperforms all models for both the *Compositional* and *Few-shot* task. Interestingly, the MF baseline is relatively strong compared to the previous methods (STRG and STIN). ORViT provides large improvement over both the previous methods and the baseline MF model. We also include results for a strong combined version (MF+STRG+STIN) of the previous methods with MF as the backbone.

#### 4.2. Spatio-temporal Action Detection

Next, we evaluate ORViT on the task of spatio-temporal action detection on the AVA dataset. In the literature, the action detection task on AVA is formulated as a two stage prediction procedure. The first step is the detection of bounding boxes, which are obtained through an off-the-shelf pre-trained person detector. The second step involves predicting the action being performed at each detected bounding box. The performance is benchmarked on the end result of these steps and is measured with the Mean Average Precision (MAP) metric. Typically, for fair comparison, the detected person boxes are kept identical across approaches and hence, the final performance depends directly on the ability of the approach to utilize the video and box information.

We follow [21, 84], using their proposed procedure and the bounding boxes they provided for both ORViT and evaluation. This enables fair comparison with all previous methods following this standard procedure.<sup>5</sup>

This task presents an ideal benchmark for evaluating the benefit of ORViT since all baselines, as well as our model, operate on the same boxes. We trained the MViT-B,  $16 \times 4$  and MViT-B,  $32 \times 3$  models on Kinetics-400 [46] and report these results. Table 2 shows that ORViT-MViT achieves +1.1, +0.7 MAP improvements over the MViT-B  $16 \times 4$  and MViT-B  $32 \times 3$ , thereby showcasing the power of our proposed object-centric representation fusion scheme.

<sup>5</sup>Boxes are obtained by pre-training FasterRCNN with a ResNeXt101-FPN [32, 59] on ImageNet and COCO human keypoint images as in [21].

(a) Something–Something V2				(b) Diving48			(c) Epic-Kitchens100					
Model	Pretrain	Top-1	Top-5	Model	Pretrain	Frames	Top-1	Method	Pretrain	A	V	N
SlowFast, R101 <sup>†</sup>	K400	63.1	87.6	SlowFast, R101 <sup>†</sup>	K400	16	77.6	SlowFast, R50 <sup>†</sup>	K400	38.5	65.6	50.0
ViViT-L <sup>†</sup>	IN+K400	65.4	89.8	TimeSformer <sup>†</sup>	IN	16	74.9	ViViT-L <sup>†</sup>	IN+K400	44.0	66.4	56.8
MViT-B, 64 <sup>†</sup>	K600	68.7	91.5	TimeSformer-L <sup>†</sup>	IN	96	81.0	MF <sup>†</sup>	IN+K400	43.1	66.7	56.5
MF <sup>†</sup>	IN+K400	66.5	90.1	TQN <sup>†</sup>	K400	ALL	81.8	MF-L <sup>†</sup>	IN+K400	44.1	67.1	57.6
MF+STRG	IN+K400	66.1	90.0	TimeSformer <sup>†</sup>	IN	32	80.0	MF-HR <sup>†</sup>	IN+K400	44.5	67.0	58.5
MF+STIN	IN+K400	66.5	89.8	TimeSformer + STRG	IN	32	78.1	MF-HR + STRG	IN+K400	42.5	65.8	55.4
MF+STRG+STIN	IN+K400	66.6	90.0	TimeSformer + STIN	IN	32	81.0	MF-HR + STIN	IN+K400	44.2	67.0	57.9
MF-L <sup>†</sup>	IN+K400	68.1	91.2	TimeSformer + STRG + STIN	IN	32	83.5	MF-HR + STRG + STIN	IN+K400	44.1	66.9	57.8
<b>ORViT MF(Ours)</b>	IN+K400	<b>67.9 (+1.4)</b>	<b>90.5 (+0.4)</b>	<b>ORViT TimeSformer(Ours)</b>	IN	32	<b>88.0 (+8.0)</b>	<b>ORViT MF-HR(Ours)</b>	IN+K400	<b>45.7 (+1.2)</b>	<b>68.4 (+1.4)</b>	<b>58.7 (+2)</b>
<b>ORViT MF-L(Ours)</b>	IN+K400	<b>69.5 (+1.4)</b>	<b>91.5 (+0.3)</b>									

Table 3. **Comparison to state-of-the-art on video action recognition.** We report top-1 (%) and top-5 (%) accuracy on SSv2. On Epic-Kitchens100 (EK100), we report top-1 (%) action (A), verb (V), and noun (N) accuracy. On Diving48 we report top-1 (%). Difference between baselines and ORViT is denoted by (+x). IN refers to IN-21K. We denote methods that do not use bounding boxes with <sup>†</sup>. For additional results and details, including the model size, see section D.2 in supplementary.

### 4.3. Action Recognition

Table 3 reports results on the standard action recognition task for several datasets. In contrast to the other tasks presented in this paper, using bounding boxes in action recognition is not part of the task definition. Thus, the comparison should be made carefully while differentiating between non-box and box-based methods. For box-based methods, we consider *STIN*, *STRG* and their combination on top of the same backbone as ORViT. Next, we explain how the boxes are extracted. For more details about datasets and evaluation see Section A of Supplementary.

**Box Input to ORViT.** For SSv2, we finetune Faster-RCNN [67] using the annotated boxes as in [63]. For EK100 and Diving48 we use Faster-RCNN [67] pre-trained on MS COCO [60]. We only use the detector bounding boxes ignoring the object classes. There is no weight sharing between the detector and our model.

**SSv2.** Table 3a shows that ORViT outperforms recent methods. The improvement is 1.4% for both *MF* and *MF-L*, while ORViT also outperforms other box-based methods, such as *MF+STIN*, *MF+STRG* and their combination. We note that these models do not improve over MF, suggesting that using boxes is non-trivial on large datasets. We also experimented with using manually annotated boxes (as opposed to those obtained by a detector) as an “oracle” upper bound to see the potential with annotated box inputs. The results for this oracle evaluation (see Section D.2 in supp) are improvements of 7.3% and 6.7% over *MF* and *MF-L* respectively. This indicates that future improvements in object detectors will benefit object-centric approaches.

**Diving48.** Here we build ORViT on top of the TimeSformer model, which was previously reported on this dataset (this demonstrates the ease of adding ORViT to any transformer model). Table 3b shows that our *ORViT TimeSformer* model outperforms the state-of-the-art methods, including TQN [97] by a significant margin of 6.2%. We obtain an improvement of 8.0% over the baseline TimeS-

former model to which ORViT blocks were added. This again indicates the direct improvement due to the ORViT block. We note that ORViT achieves these results using only 32 frames, significantly less than the previous best results, *TimeSformer-L*, which uses 96 frames. ORViT outperforms box-based methods, including *TimeSformer+STIN+STRG* (4.5%), *TimeSformer+STIN* (7.0%), and *TimeSformer+STRG* (9.9%).

**EK100.** Table 3c reports results on EK100. Here we add ORViT blocks to the *MF-HR* model (which is the best performing model on EK100 in [65]). Results show that our *ORViT MF-HR* model improves the accuracy for all three sub-tasks (with a smaller improvement for nouns). We believe the improvements on EK100 are less impressive than on the other datasets for two main reasons: (a) EK100 is an ego-centric dataset, making the camera movement a significant challenge for our method to model meaningful object interactions. (b) EK100 contains short 2-3 seconds videos, thus temporal reasoning is less effective.

### 4.4. Ablations

We perform a comprehensive ablation study on the compositional action recognition task [63] in the “SomethingElse” dataset to test the contribution of the different ORViT components (Table 4). We use the MF as the baseline architecture for ORViT. For more ablations, see Section C in supplementary.

**Components of the ORViT model.** We consider the following versions of our model. (i) Single ORViT block (no ODM stream<sup>6</sup>). We first consider a single application of the ORViT block, but without the ODM stream. We also compare different video transformer layers at which to apply our ORViT block (namely, the video transformer layer from which to extract the RoI descriptors). We refer to models applied at layer  $X$  as *ORViT[L:X]*. (ii) Single ORViT block (with ODM stream). Here we augment the single ORViT

<sup>6</sup>We refer to the “Object-Dynamics Module” as ODM stream.

(a) Components					(b) Object-centric Baselines			(c) Boxes			(d) ODM Dimension				
Layers	Top-1	Top-5	GFLOP	Param	Model	Top-1	Top-5	Model	Top-1	Top-5	Dim.	Top-1	Top-5	GFLOP	Param
MF	60.2	85.8	$\times 1(370)$	$\times 1(109)$	MF	60.2	85.8	Full boxes	60.9	84.5	MF	60.2	85.8	$\times 1(370)$	$\times 1(109)$
ORViT [L:12]	63.9	87.6	$\times 1.01$	$\times 1.01$	MF + RoIAlign	59.6	84.5	Null boxes	60.4	84.2	0	67.4	89.8	$\times 1.03$	$\times 1.02$
ORViT [L:2]	66.7	89.2	$\times 1.01$	$\times 1.01$	MF + Boxes	63.7	86.9	Grid boxes	60.9	84.8	128	68.7	90.3	$\times 1.03$	$\times 1.03$
ORViT [L:2]+ODM	68.8	90.5	$\times 1.03$	$\times 1.12$	<b>ORViT (Ours)</b>	<b>69.7</b>	<b>91.0</b>	Random boxes	60.7	85.0	256	68.9	90.5	$\times 1.03$	$\times 1.05$
ORViT [L:2,7,11]+ODM	<b>69.7</b>	<b>91.0</b>	$\times 1.09$	$\times 1.36$				<b>Object Regions (Ours)</b>	<b>69.7</b>	<b>91.0</b>	768	69.7	91.0	$\times 1.1$	$\times 1.36$

Table 4. **Ablications.** We report top-1 and top-5 action accuracy on the SomethingElse split. We show (a) Contribution of ORViT components (with parameters number in  $10^6$  and GFLOPS in  $10^9$ ). (b) Other Object-centric baselines. (c) ORViT with different boxes input, and (d) The effect of ‘‘Object-Dynamics Module’’ (ODM) embedding dimension. More ablations are in Section C in supplementary.

block, with the ODM stream. We refer to these models as *ORViT[L:X]+ODM*. (iii) Multiple ORViT blocks (with ODM stream). This is the version of ORViT used in all our experiments. It applies the ORViT block at multiple layers. We chose layers 2,7 and 11 of the video transformer model to apply ORViT block at. All the ablations were performed on the compositional split in SomethingElse. In the ablation table, we refer to this as *ORViT[L:2,7,11]+ODM*. In the rest of the experiments this is simply referred to as ORViT.

The results are shown in Table 4a. It can be seen that a single ORViT layer already results in considerable improvement (66.7%), and that it is very important to apply it in the earlier layers rather than at the end. This is in contrast to the current practice in object-centric approaches (e.g., STRG and STIN) that extract RoIs from the final layer. It can also be seen that the ODM stream improves performance (by 2.1% from 66.7% to 68.8%). Finally, multiple applications of the layer further improve performance to 69.7%.

**Object-Centric Baselines.** ORViT proposes an elegant way to integrate *object region information* into a video transformer. Here we consider two other candidate models to achieve this goal. (i) *MF+RoIAlign* uses RoIAlign over the last video transformer layer to extract object features. Then, it concatenates the CLS token with max-pooled object features to classify the action using an MLP. (ii) *MF+Boxes* uses coordinates and patch tokens. We use the CLS token from the last layer of MF, concatenated with trajectory embeddings. To obtain trajectory embeddings, we use a standard self-attention over the coordinates similar to our ODM stream. The first captures the appearance of objects with global context while the latter captures the trajectory information with global context, both without fusing the object information several times back to the backbone as we do. The results are shown in Table 4b. *MF+RoIAlign* does not improve over the baseline, while *MF+Boxes* improves by 3.5%, which is still far from ORViT (69.7%).

**How important are the object bounding boxes.** Since ORViT changes the architecture of the base video transformer model, we want to check whether the bounding boxes are indeed the source of improvement. We consider several variations where the object bounding boxes are replaced with other values. (i) *All boxes*: all boxes are given the coordinates of the entire image  $([0, 0, 1, 1])$ . (ii) *Null*

*boxes*: boxes are initialized to zeros. (iii) *Grid boxes*: each of the 4 bounding boxes is one fourth of the image. (iv) *Random boxes* - each box is chosen uniformly at random. See Table 4c for results. We observe a large drop in performance for all these baselines, which confirms the important role of the object regions in ORViT. Finally, we ask whether tracking information is important, as opposed to just detection. We find that this results in degradation from 69.7 to 68.2, indicating that the model can perform relatively well with only detection information.

**Decreasing Model Size.** Next, we show that model size can be significantly decreased, incurring a small performance loss. Most the parameters added by ORViT over the baseline MF are in the ODM, and thus it is possible to use a smaller embedding dimension in ODM (see  $\tilde{B}$  in Section 3.2). Table 4d reports how the dimension affects the performance, demonstrating that most of the performance gains can be achieved with a model that is close in size to the original MF. More in D.1&D.2 in supp. We would like to highlight that ‘‘Object-Region Attention’’ alone (set dimension size to 0; thus ODM is not used) is the main reason for the improvement with only 2% additional parameters.

## 5. Discussion and Limitations

Objects are a key element of human visual perception, but their modeling is still a challenge for machine vision. In this work, we demonstrated the value of an object-centric approach that incorporates object representations starting from early layers and propagates them into the transformer-layers. Through extensive empirical study, we show that integrating the ORViT block into video transformer architecture leads to improved results on four video understanding tasks and five datasets. However, we did not put effort into the object detection and used externally provided boxes, which is a limitation of our work. Replacing the externally provided boxes with boxes that the model generates without supervision will be interesting.

**Acknowledgements.** This project has received funding from the ERC under the European Unions Horizon 2020 research and innovation programme (grant ERC HOLI 819080). Prof. Darrell’s group was supported in part by DoD including DARPA’s XAI, and LwLL programs, as well as BAIR’s industrial programs. This work was completed in partial fulfillment for the Ph.D. degree of the first author.



## References

- [1] Humam Alwassel, Fabian Caba Heilbron, and Bernard Ghanem. Action search: Spotting targets in videos and its application to temporal action localization. In *The European Conference on Computer Vision (ECCV)*, September 2018. [3](#)
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer, 2021. [2](#), [3](#), [16](#)
- [3] Anurag Arnab, Chen Sun, and Cordelia Schmid. Unified graph structured models for video understanding. In *ICCV*, 2021. [2](#)
- [4] Amir Bar, Roei Herzig, Xiaolong Wang, Anna Rohrbach, Gal Chechik, Trevor Darrell, and A. Globerson. Compositional video synthesis with action graphs. In *ICML*, 2021. [2](#)
- [5] Fabien Baradel, Natalia Neverova, Christian Wolf, Julien Mille, and Greg Mori. Object level visual reasoning in videos. In *ECCV*, pages 105–121, 2018. [2](#)
- [6] 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. *arXiv preprint arXiv:1806.01261*, 2018. [2](#)
- [7] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *Proceedings of the International Conference on Machine Learning (ICML)*, July 2021. [2](#), [3](#), [5](#), [6](#), [13](#), [14](#), [16](#)
- [8] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 3464–3468, 2016. [6](#), [13](#)
- [9] S. Buch, V. Escorcia, C. Shen, B. Ghanem, and J. C. Niebles. Sst: Single-stream temporal action proposals. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6373–6382, 2017. [3](#)
- [10] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 213–229, 2020. [2](#)
- [11] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. [3](#), [6](#)
- [12] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A. Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster R-CNN architecture for temporal action localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. [3](#)
- [13] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. [2](#)
- [14] Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. Randaugment: Practical automated data augmentation with a reduced search space. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 18613–18624. Curran Associates, Inc., 2020. [13](#), [14](#)
- [15] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S. Davis, and Yan Qiu Chen. Temporal context network for activity localization in videos. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [3](#)
- [16] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, , Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision. *CoRR*, abs/2006.13256, 2020. [2](#), [5](#), [13](#)
- [17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. [2](#)
- [18] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor 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. [3](#)
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. [2](#), [3](#)
- [20] Alexei A Efros, Alexander C Berg, Greg Mori, and Jitendra Malik. Recognizing action at a distance. In *Computer Vision, IEEE International Conference on*, volume 3, pages 726–726. IEEE Computer Society, 2003. [3](#)
- [21] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6201–6210, 2019. [3](#), [6](#), [13](#), [14](#)
- [22] Christoph Feichtenhofer, Axel Pinz, and Andrew 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. [3](#)
- [23] Chen Gao, Jiarui Xu, Yuliang Zou, and Jia-Bin Huang. Drg: Dual relation graph for human-object interaction detection. *ArXiv*, abs/2008.11714, 2020. [2](#)
- [24] Jiyang Gao, Zhenheng Yang, Kan Chen, Chen Sun, and Ram Nevatia. Turn tap: Temporal unit regression network for temporal action proposals. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [3](#)
- [25] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *Proceed-*

- ings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 244–253, 2019. [2](#)
- [26] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, page 5, 2017. [2](#), [5](#), [13](#)
- [27] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6047–6056, 2018. [3](#)
- [28] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*, pages 6047–6056. IEEE Computer Society, 2018. [2](#), [5](#)
- [29] Abhinav Gupta and Larry S. Davis. Objects in action: An approach for combining action understanding and object perception. *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. [2](#)
- [30] Fan Haoqi, Xiong Bo, Mangalam Karttikeya, Li Yanghao, Yan Zhicheng, Malik Jitendra, and Feichtenhofer Christoph. Multiscale vision transformers. *arXiv:2104.11227*, 2021. [2](#), [3](#), [5](#), [6](#), [13](#), [14](#), [16](#)
- [31] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017. [4](#), [13](#), [15](#)
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [6](#), [13](#)
- [33] Roee Herzig, Amir Bar, Huijuan Xu, Gal Chechik, Trevor Darrell, and Amir Globerson. Learning canonical representations for scene graph to image generation. In *European Conference on Computer Vision*, 2020. [2](#)
- [34] Roee Herzig, Elad Levi, Huijuan Xu, Hang Gao, Eli Brosh, Xiaolong Wang, Amir Globerson, and Trevor Darrell. Spatio-temporal action graph networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. [2](#)
- [35] Roee Herzig, Moshiko Raboh, Gal Chechik, Jonathan Berant, and Amir Globerson. Mapping images to scene graphs with permutation-invariant structured prediction. In *Advances in Neural Information Processing Systems (NIPS)*, 2018. [2](#)
- [36] Geoffrey E. Hinton, Nitish Srivastava, A. Krizhevsky, Ilya Sutskever, and R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *ArXiv*, abs/1207.0580, 2012. [13](#), [14](#)
- [37] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In *ECCV*, 2016. [13](#), [14](#)
- [38] Max Jaderberg, K. Simonyan, Andrew Zisserman, and K. Kavukcuoglu. Spatial transformer networks. In *NIPS*, 2015. [5](#)
- [39] M. Jain, A. Ghodrati, and C. G. M. Snoek. Actionbytes: Learning from trimmed videos to localize actions. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1168–1177, 2020. [3](#)
- [40] Achiya Jerbi, Roee Herzig, Jonathan Berant, Gal Chechik, and Amir Globerson. Learning object detection from captions via textual scene attributes. *ArXiv*, abs/2009.14558, 2020. [2](#)
- [41] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):221–231, 2013. [3](#)
- [42] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image generation from scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1219–1228, 2018. [2](#), [5](#)
- [43] R. Kálmán. A new approach to linear filtering and prediction problems” transaction of the asme journal of basic. In *transaction of the asme journal of basic*, 1960. [6](#), [13](#)
- [44] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li 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. [3](#)
- [45] Keizo Kato, Yin Li, and Abhinav Gupta. Compositional learning for human object interaction. In *ECCV*, 2018. [2](#)
- [46] 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. *arXiv preprint arXiv:1705.06950*, 2017. [6](#), [13](#), [14](#)
- [47] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [13](#), [14](#)
- [48] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. [16](#)
- [49] Ranjay Krishna, Ines Chami, Michael S. Bernstein, and Li Fei-Fei. Referring relationships. *ECCV*, 2018. [2](#)
- [50] H. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955. [13](#)
- [51] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager. Temporal convolutional networks for action segmentation and detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1003–1012, 2017. [3](#)
- [52] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *ArXiv*, abs/1908.03557, 2019. [2](#)

- [53] Xiujun Li, Xi Yin, Chunyuan Li, Xiaowei Hu, Pengchuan Zhang, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantic aligned pre-training for vision-language tasks. *ECCV 2020*, 2020. [2](#)
- [54] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [2](#), [6](#), [14](#)
- [55] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G. Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5718–5727, 2019. [2](#)
- [56] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [3](#)
- [57] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. BMN: boundary-matching network for temporal action proposal generation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV*, pages 3888–3897. IEEE, 2019. [3](#)
- [58] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, 2018. [3](#)
- [59] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 936–944, 2017. [6](#), [13](#)
- [60] Tsung-Yi Lin, M. Maire, Serge J. Belongie, James Hays, P. Perona, D. Ramanan, Piotr Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [7](#), [13](#)
- [61] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 344–353. Computer Vision Foundation / IEEE, 2019. [3](#)
- [62] S. Ma, L. Sigal, and S. Sclaroff. Learning activity progression in lstms for activity detection and early detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1942–1950, 2016. [3](#)
- [63] Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, and Trevor Darrell. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. [2](#), [3](#), [5](#), [6](#), [7](#), [13](#), [14](#), [15](#)
- [64] Megha Nawhal and Greg Mori. Activity graph transformer for temporal action localization, 2021. [2](#)
- [65] Mandela Patrick, Dylan Campbell, Yuki M. Asano, Ishan Misra Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and Joao F. Henriques. Keeping your eye on the ball: Trajectory attention in video transformers, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [13](#), [14](#), [15](#), [16](#)
- [66] Moshiko Raboh, Roei Herzig, Gal Chechik, Jonathan Berant, and Amir Globerson. Differentiable scene graphs. In *WACV*, 2020. [2](#)
- [67] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. [7](#), [13](#)
- [68] A. Richard and J. Gall. Temporal action detection using a statistical language model. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3131–3140, 2016. [3](#)
- [69] Kate Saenko, Ben Packer, Chao-Yeh Chen, Sunil Bandla, Yong Jae Lee, Yangqing Jia, Juan Carlos Niebles, Daphne Koller, Li Fei-Fei, Kristen Grauman, and Trevor Darrell. Mid-level features improve recognition of interactive activities. *arXiv preprint arXiv:1912.06992*, 2012. [2](#)
- [70] 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. [2](#)
- [71] Baifeng Shi, Qi Dai, Yadong Mu, and Jingdong Wang. Weakly-supervised action localization by generative attention modeling. *arXiv preprint arXiv:2003.12424*, 2020. [3](#)
- [72] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pages 1049–1058. IEEE Computer Society, 2016. [3](#)
- [73] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014. [3](#)
- [74] Bharat Singh, Tim K. Marks, Michael Jones, Oncel Tuzel, and Ming Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [3](#)
- [75] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 318–334, 2018. [2](#)
- [76] Xu Sun, Tongwei Ren, Yuan Zi, and Gangshan Wu. Video visual relation detection via multi-modal feature fusion. *Proceedings of the 27th ACM International Conference on Multimedia*, 2019. [2](#)
- [77] Hao Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *EMNLP*, 2019. [2](#)

- [78] Graham W Taylor, Rob Fergus, Yann LeCun, and Christoph Bregler. Convolutional learning of spatio-temporal features. In *European conference on computer vision*, pages 140–153. Springer, 2010. 3
- [79] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 3
- [80] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 3
- [81] Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1510–1517, 2018. 3
- [82] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(11):2740–2755, 2019. 3
- [83] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018. 2, 6
- [84] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krähenbühl, and Ross Girshick. Long-Term Feature Banks for Detailed Video Understanding. In *CVPR*, 2019. 6, 13
- [85] Chao-Yuan Wu and Philipp Krähenbühl. Towards Long-Form Video Understanding. In *CVPR*, 2021. 3
- [86] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 13
- [87] Bingjie Xu, Yongkang Wong, Junnan Li, Qi Zhao, and M. Kankanhalli. Learning to detect human-object interactions with knowledge. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2019–2028, 2019. 2
- [88] H. Xu, A. Das, and K. Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5794–5803, 2017. 3
- [89] Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S. Du, Ken-ichi Kawarabayashi, and Stefanie Jegelka. What can neural networks reason about? In *International Conference on Learning Representations*, 2020. 2
- [90] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 3
- [91] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5812–5823. Curran Associates, Inc., 2020. 2
- [92] Yuan Yuan, Yueming Lyu, Xi Shen, Ivor W. Tsang, and Dit-Yan Yeung. Marginalized average attentional network for weakly-supervised learning. In *International Conference on Learning Representations*, 2019. 3
- [93] Zehuan Yuan, Jonathan C. Stroud, Tong Lu, and Jia Deng. Temporal action localization by structured maximal sums. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 3
- [94] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George 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. 3
- [95] Vinicius Zambaldi, David Raposo, Adam Santoro, Victor Bapst, Yujia Li, Igor Babuschkin, Karl Tuyls, David Reichert, Timothy Lillicrap, Edward Lockhart, et al. Relational deep reinforcement learning. *arXiv preprint arXiv:1806.01830*, 2018. 2
- [96] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *ICCV*, 2019. 3
- [97] Chuhan Zhang, Ankush Gputa, and Andrew Zisserman. Temporal query networks for fine-grained video understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 7, 14
- [98] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, 2017. 3
- [99] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. *European Conference on Computer Vision*, 2018. 2
- [100] Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cinbis, David Fouhey, Ivan Laptev, and Josef Sivic. Cross-task weakly supervised learning from instructional videos. In *CVPR*, 2019. 3