

# Action-slot: Visual Action-centric Representations for Multi-label Atomic Activity Recognition in Traffic Scenes

Chi-Hsi Kung<sup>1</sup>

Shu-Wei Lu<sup>1</sup>

Yi-Hsuan Tsai<sup>2</sup>


Yi-Ting Chen<sup>1</sup>

National Yang Ming Chiao Tung University<sup>1</sup>

{chkung, tomy45651.sc06, ychen}@nycu.edu.tw

Google<sup>2</sup>

yhtsai@google.com



**Figure 1.** Illustration of the concept of multi-label atomic activity recognition and our proposed Action-slot. In the scene, three atomic activities are presented and depicted by colored arrows. For example, the red arrow represents the **Z1-Z4: C+** atomic activity, indicating a group of vehicles turning left. Atomic activities are defined based on road user’s type and their motion patterns grounded in the underlying road structure. We introduce Action-slot to learn visual action-centric representations that enable decomposing multiple atomic activities in videos. We demonstrate that our framework can effectively recognize multiple atomic activities via learned representations.

## Abstract


In this paper, we study multi-label atomic activity recognition. Despite the notable progress in action recognition, it is still challenging to recognize atomic activities due to a deficiency in holistic understanding of both multiple road users’ motions and their contextual information. In this paper, we introduce Action-slot, a slot attention-based approach that learns visual action-centric representations, capturing both motion and contextual information. Our key idea is to design action slots that are capable of paying attention to regions where atomic activities occur, without the need for explicit perception guidance. To further enhance slot attention, we introduce a background slot that competes with action slots, aiding the training process in avoiding unnecessary focus on background regions devoid of activities. Yet, the imbalanced class distribution in the existing dataset hampers the assessment of rare activities. To address the limitation, we collect a synthetic dataset called TACO, which is four times larger than OATS and features a balanced distribution of atomic activities. To validate the effectiveness of our method, we conduct comprehensive experi-

ments and ablation studies against various action recognition baselines. We also show that the performance of multi-label atomic activity recognition on real-world datasets can be improved by pretraining representations on TACO. Our source code, dataset, and visualization videos are available at <https://hcis-lab.github.io/Action-slot/>

## 1. Introduction

There has been a noteworthy rise of community interest in activity recognition within traffic scenes [3, 13, 40, 43, 46, 47, 53]. Recognizing activities of road users is essential for advancing the development of intelligent driving systems, as it enables various applications, such as intent prediction [30, 40, 61], scenario retrieval [39, 45, 51, 55], scenario-based assessment [1, 62], and safety-critical scenario generation [10, 48, 49, 62].

In this paper, we focus on multi-label atomic activity recognition, a new task proposed by the OATS dataset for interactive traffic scenario understanding [3]. An atomic activity is a higher-level semantic motion pattern rooted in the



**Figure 2.** The distribution of atomic activity classes in our TACO dataset compared to the OATS dataset. Please note that, for space considerations, we omit the road topology notations of corner (C) and roadway (Z) on the x-axis of the figure.

underlying road topology. For instance, a group of vehicles turning left is described as **Z1-Z4: C+**, as shown in Figure 1. The notions *Z1* and *Z4* are roadways within an intersection. The left-turn motion pattern is represented as **Z1-Z4**. We use *C+* to denote a group of vehicles. The combination of such motion patterns and road user types can compose 64 classes of atomic activities. This task presents distinctive challenges as it demands a comprehensive understanding of contextual information and motions of road users. Moreover, it requires the ability to break down activities within crowded and intricate traffic scenes from videos.

We identify the limitations of current action recognition models, showing less desirable performance in the OATS benchmark [3]. We argue that video-level representations, simply employing 3D ConvNet [12, 23, 24, 57] or transformers [22, 56] models, pose a challenge for classifiers in distinguishing individual atomic activities within traffic scenes due to the complicated structure of traffic scenes. On the other hand, object-aware representations [3, 6, 60] using object features [28] struggle to associate the relationship between the road structures and motions of objects. Then the question we like to answer in this paper: *Can we learn visual representations that decompose multiple atomic activities from videos, without using object proposals?*

To this end, we propose Action-slot, a slot attention-based framework, inspired by the recent success of slot attention for unsupervised object discovery [41]. Action-slot incorporates four crucial design choices that allow the model to decompose multiple atomic activities from videos and represent them without relying on object proposals. First, we allocate a fixed number of action slots, assigning each action slot to focus on regions where a specific atomic activity occurs, such as **C2-C1: P**. Second, we introduce an additional background slot and attention guidance to enforce the slot paying attention to regions without active atomic activities. Third, we discourage action slots allocated with negative classes (i.e., those activities not present in a video), from attending to any regions. Fourth, we modify the slot updating strategy used in [5, 20, 34] to a parallel updating approach that integrates spatial-temporal information at the video level into a slot.

In our experiments, we demonstrate the effectiveness of

*Action-slot* in accurately identifying a variety of atomic activities annotated in the OATS dataset [3]. Furthermore, we conduct thorough experiments, comparing various action recognition baselines and presenting ablation studies that validate the design of Action-slot. Moreover, we provide qualitative evidence demonstrating that the learned action-centric representations can reliably identify distinct atomic activities in the visual domain, even without the need for supervision signals such as object locations [20]. However, due to imbalanced class distribution within the dataset, OATS only uses 35 out of 64 atomic activity classes for training and evaluation. Specifically, there are no activities associated with group two-wheelers (*K+*). This observation extends to nearly half of the activity classes, impeding the evaluation of rare classes.

To address this limitation, we introduce the Traffic Activity Recognition (TACO) dataset, an extensive dataset for atomic activity recognition. We utilize the CARLA simulator [18] to gather instances of all conceivable activity classes, ensuring a well-balanced distribution, as illustrated in Figure 2. We again benchmark all methods on TACO and show that Action-slot outperforms all baselines by a large margin.

Beyond comprehensive performance analysis, we demonstrate that TACO can enhance the efficacy of atomic activity recognition on real-world datasets through pretraining feature representations on TACO. Then we finetune the pretrained model on both the OATS dataset and a newly annotated nuScenes dataset [9]. We present significant improvements in all experiments through TACO pre-training, underscoring the robust real-world transferability of TACO. Our main contributions are as follows:

1. We introduce *Action-slot*: an action-centric slot attention-based framework that decomposes multiple atomic activities from videos.
2. We present the TACO dataset, an extensive and balanced dataset for multi-label atomic activity recognition, facilitating comprehensive performance analysis.
3. We conduct extensive evaluations and ablation studies across multiple datasets to justify the effectiveness and generalization of Action-slot.
4. We provide qualitative evidence demonstrating that the

learned action-centric representations can reliably identify distinct atomic activities in the visual domain without perception modules.

## 2. Related Work

**Traffic Scenario Understanding Datasets.** The development of intelligent driving systems, such as scene analysis [40, 46, 61], scenario retrieval [39, 45, 51, 55], and safety-critical scenario generation [10, 48, 49, 62], heavily relies on traffic activity understanding. While previous efforts have designed datasets [13, 38, 43, 46] by labeling high-level actions (e.g., left turn) for traffic pattern recognition, they fall short in supporting fine-grained scenario analysis. For instance, high-level actions cannot differentiate “left turns” that start from different locations, such as the ego lane versus the oncoming lane. To address this limitation, both Inner-City [13] and ROAD [53] datasets augment high-level action labels (e.g., turning left) with location labels (e.g., on the oncoming lane or on our side of the road). Recently, OATS [3] propose a topology-aware traffic activity description language to unify the action and location labels via the topology-aware pattern description. Specifically, the language decomposes an interactive scenario into a set of atomic activities, that are defined based on the types of road users and the corresponding motion patterns grounded in road structures. Nevertheless, some real-world instances of traffic atomic activities are scarce, presenting a challenge in forming a well-balanced category distribution and impeding the evaluation of rare classes. In response, we construct the TACO dataset utilizing the CARLA simulator [18] to gather instances of all conceivable activity classes, ensuring a well-balanced distribution.

**Video Action Recognition.** Substantial improvements in spatial-temporal modeling via 3D ConvNet [12, 23, 24, 57] and transformer [4, 7, 22, 56] have been observed because of the Kinetics human action dataset [33] that facilitates model pre-training. The pre-trained models can further be finetuned on other action recognition datasets such as the HMDB [36], THUMOS [31], and Charades [52] datasets. Meanwhile, the community also explores multi-label action recognition and establishes benchmarking datasets, such as MultiTHUMOS [64], Charades [52], and AVA [26]. The most relevant task to multi-label atomic activity recognition is group activity recognition [16, 29], where multiple players engage in distinct actions. In contrast, atomic activity recognition in traffic scenes demands a comprehensive understanding of objects and their motion patterns within the underlying topology. In addition, a noteworthy distinction is that the majority of objects in traffic scenes do not exhibit any activity, constituting a negative class, a concept absent in group activity recognition datasets.

**Slot Attention for Representation Learning.** Locatello et al., [41] propose slot attention as a new perspective for learning object-centric representations in an unsupervised manner. They have successfully stimulated the community to explore strategies to bootstrap slot attention from synthetic datasets [32] to real-world dynamic visual scenes with moving cameras such as the Waymo [54] and KITTI [25] datasets. Particularly, they utilize auxiliary information such as optical flow [34], motion segmentation [5], object initial locations [20], and depth [20] to facilitate object-centric representation learning in challenging real-world applications.

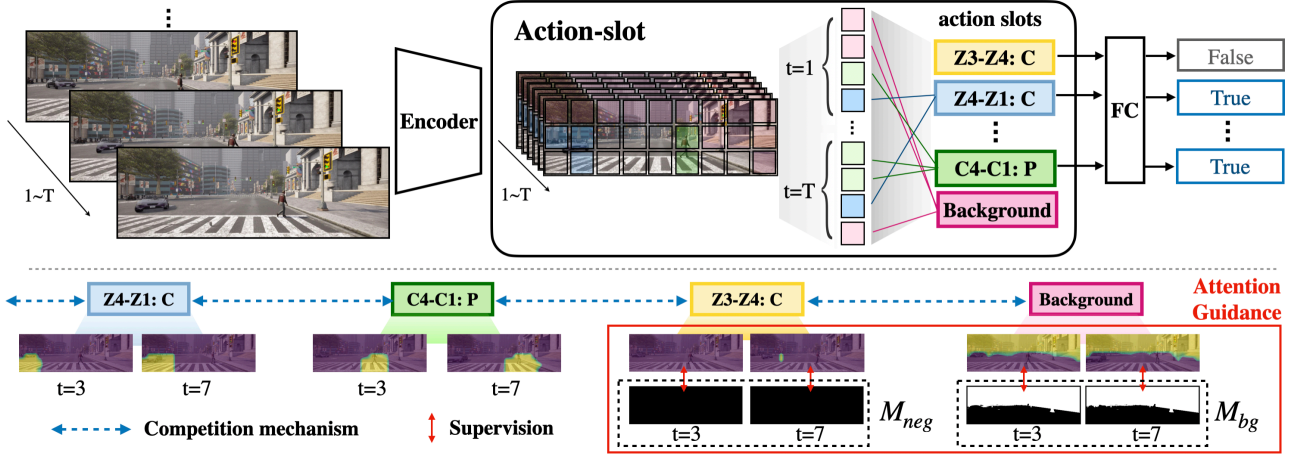
Multi-label atomic activity recognition poses three challenges for existing slot attention methods. First, the slots are permutation invariant. This property is designed to discover arbitrary objects [5, 8, 20, 34, 41, 50, 65] but is not suitable for a classification task with fixed-length classes. Second, current slot attention architectures used in videos, such as SAVi [34] and Slot-VPS [65], are designed for object-centric tasks, such as tracking. Specifically, they update a slot for each frame in a recurrent manner, which is not suitable for action recognition that requires a holistic temporal understanding. Third, to enable slot attention in complex scenes, they leverage additional object signals to guide the attention during training, such as flow [34], depth [20], or motion segmentation [5]. Object signals may lead to misinterpretations in atomic activity recognition since not all objects participate in activities (e.g., pedestrians strolling on the sidewalk and vehicles waiting at traffic lights). In this work, we propose a novel action-centric slot attention-based framework and demonstrate the effectiveness of the framework on multiple datasets.

## 3. Method

In this section, we first introduce the background of slot attention. Then, we provide an overview of the problem definition in multi-label atomic activity recognition and the proposed method. Finally, we explain the specific modifications we make to the slot attention for our task.

### 3.1. Preliminary: Slot Attention

The Slot Attention module [41] can be viewed as a clustering algorithm that maps image patches input to a set of  $K$  output slots  $S \in \mathbb{R}^{K \times D_{\text{slot}}}$ , where  $D_{\text{slot}}$  is the dimension of each slot. Specifically, the slots  $S$  are first initialized by randomly sampling  $K$  vectors from a Gaussian distribution, where the parameter  $K$  is usually defined as the maximum number of objects in an image. An input frame, i.e., image features  $F \in \mathbb{R}^{H \times W \times D_{\text{in}}}$  with size  $H \times W$  and dimension  $D_{\text{in}}$ , is first flattened to patch tokens  $F' \in \mathbb{R}^{N \times D_{\text{in}}}$ , where  $N = H \times W$ . Then the tokens are mapped to the slots using the dot product attention module. That is, the attention weight can be calculated as  $A = \frac{1}{\sqrt{D}} k(F') \cdot q(S) \in \mathbb{R}^{N \times K}$ ,



**Figure 3.** The top of the figure illustrates the proposed framework. Action-slot takes video as input and uses a CNN encoder to extract feature patches. All patches are then processed with individual slots simultaneously to find the most relevant spatial-temporal patches corresponding to each action slot. The updated action slots are fed into a fully connected layer to predict the probability of the corresponding action class, excluding the background slot. The bottom of the figure depicts the attention maps of action and background slots. We propose to incorporate a background mask  $M_{bg}$  to supervise the background slot. The design facilitates other action slots to capture action signals. Furthermore, we design a regularization for slots allocated to negative classes (e.g., Z3-Z4: C) using an all-zero mask  $M_{neg}$ .

and  $q$  and  $k$  are linear transformations that map the input  $F$  and slots  $S$  to a common dimension  $D$ .

We obtain the updated values for slots through  $U = \bar{A}^T \cdot v(F^t) \in \mathbb{R}^{K \times D}$ , where  $\bar{A}$  is the normalized attention weight calculated via the softmax operation and  $v$  is a linear transformation. Finally, the slots are updated via a GRU:  $S^t = \text{GRU}(S, U)$  [15]. To refine the slots, the updating process repeats  $M$  iterations [20, 34, 41]. Note that, the key difference between classical attention [58] and slot attention is that the attention weights  $A$  are normalized with the softmax operation slot-wise instead of token-wise. This difference enables slots to compete with each other so that each slot attends to different relevant regions of input. To extend slot attention to video tasks, the previous works [5, 20, 34] propagate slots recurrently, i.e.,  $S^t = \text{GRU}(S^{t-1}, U)$ .

### 3.2. Overview of Proposed Method

**Problem Formulation.** Given a video clip  $V_i$  with  $T$  image frames  $\{I_t^i\}_{t=1}^T$ , our goal is to recognize whether there are any atomic activities  $Y$  presenting in the video. The ground truth of  $Y^i$  for video  $V_i$  is a binary multi-label vector, i.e.,  $Y^i = \{y_c\}_{c=1}^{N_{c1}}$ , where  $y_c$  is 1 if the corresponding activity appears and vice versa, and  $N_{c1}$  denotes the total number of possible activities. Note that, we do not constrain when a certain activity appears, e.g.,  $y_c$  can happen in any image frames of  $\{I_t^i\}_{t=1}^T$ , or may appear multiple times in one video.


**Overview of Action-slot.** We propose Action-slot, an action-centric slot-attention-based framework to decompose multiple atomic activities from videos. Our idea is to assign each slot to learn specific action-centric representations for the corresponding atomic activity. Figure 3

presents the overview of Action-slot. Different from existing slot attention algorithms designed for unsupervised object discovery and tracking [5, 20, 34], we make the following three modifications. First, we allocate a fixed number of slots  $K$ , which is equal to the number of atomic activities  $N_{c1}$ . We supervise the prediction of each action slot with the corresponding ground truth label  $y_c$ , i.e., whether the atomic activity  $c$  appears. Second, we introduce a background slot to focus on the regions that are not possible to present an activity, so that the other action slots can compete with this background slot. In other words, the design forces other action slots to pay attention to the important regions. Third, we discourage action slots associated with negative classes (i.e., those activities not present in a video), from attending to any regions. Lastly, we update all action slots in a parallel fashion by considering all the image frames together, instead of recurrently updating slots along the temporal domain utilized in [5, 20, 34].

### 3.3. Action-slot

In this section, we describe the details of our design for learning action-centric representations. We follow [5] to use learnable parameters for slot initializations. We extract image features  $F \in \mathbb{R}^{T \times H \times W \times D_{in}}$  from the convolution-based encoder [12, 23, 24, 27] for each frame in a video clip and then flatten them to  $F^t \in \mathbb{R}^{N \times D_{in}}$ , where  $N = T \times H \times W$ . We add learnable 3D spatial-temporal positional embeddings  $E \in \mathbb{R}^{T \times H \times W}$  to the tokens  $F^t$  [4, 5, 7, 22, 56]. **Allocated slot.** We then define a set of slots  $\{S_k\}_{k=1}^K$ , where  $K = N_{c1}$ , representing the number of action slots (i.e., number of activity classes). In contrast to the previous slot attention works [20, 34, 41, 65] where they do not spec-





**Figure 4.** (a) Our parallel scheme considers updating slots based on the spatial-temporal features for all the frames. (b) The slots are updated recurrently along the temporal dimension, where each time the slot only considers a frame-wise feature.

ify a task to each slot and use a Hungarian matcher for bipartite matching with objects, we instead allocate each slot a class of activity, namely, action slots. Each action slot is processed with an independent binary classifier to output a prediction  $\hat{y}_c$  for each class  $c$ , where a binary cross entropy (BCE) function is used as the objective for each data sample:  $L_{\text{act}} = \sum_{c=1}^{N_{c1}} \text{BCE}(\hat{y}_c, y_c)$ .

**Parallel updating.** We update slots across image frames in a parallel manner, as shown in Figure 4 (a). Specifically, we calculate normalized attention weight  $\bar{A} \in \mathbb{R}^{N \times K}$  while considering temporal dimension at the same time, i.e.,  $N = T \times H \times W$ . We then update the  $K$  slots with attention weights  $\bar{A}$  in a single pass. This is different from the previous works [5, 20, 34] where they update slots across frames recurrently, namely, using the previous state of slots and current feature maps to update the current state of slots for  $T$  iterations, as shown in Figure 4 (b).

**Background slot and attention guidance.** It is non-trivial to provide guidance for positive classes as the objects are not necessarily involved in activities all the time, and not all objects are involved in activities. Therefore, we propose to use a background slot that is not supervised by any activity classes. By the softmax operation in slot attention, this background slot would attend regions that are not relevant to any activities and force other action slots to focus on other regions, as shown in the bottom of Figure 3. To further enhance the effectiveness of the background slot, we provide an attention mask  $M_{\text{bg}}$  as the supervision, directly guiding the background slot where to pay attention. The mask is generated by excluding pixels with the classes of vehicles, pedestrians, drivable areas, crosswalks, and sidewalks. We guide the background attention via the loss function:  $L_{\text{bg}} = \text{BCE}(\bar{A}_{\text{bg}}, M_{\text{bg}})$ , where  $\bar{A}_{\text{bg}}$  is the normalized attention from the background slot.

**Regularization for action slots.** Instead of providing attention guidance for positive classes, which is non-trivial as explained previously, we introduce a regularization term to discourage action slots associated with negative classes (i.e., those activities not present in a video), from attending to any regions. We construct an attention mask  $M_{\text{neg}}$  wherein all element values are set to 0. We design a loss function:  $L_{\text{neg}} = \sum_{\{c|y_c=0\}} \text{BCE}(\bar{A}_c, M_{\text{neg}})$ , where  $\bar{A}_c$  is

the normalized attention output of slot  $S_c$  for negative class  $c$ . By regularizing the action slots with negative classes, the competition mechanism makes slots with positive classes more competitive in the important regions. Note that the regularization does not require additional annotations. Finally, the full objective function of Action-slot is:

$$L_{\text{all}} = L_{\text{act}} + w_{\text{bg}}L_{\text{bg}} + w_{\text{neg}}L_{\text{neg}}$$

, where  $w_{\text{bg}}$  and  $w_{\text{neg}}$  are the weights to balance two terms (set as 0.5 and 1, respectively in this paper).

## 4. Experiments

### 4.1. Datasets

**OATS [3].** The dataset was collected in San Francisco with an instrumented vehicle and comprises 1026 labeled clips and is divided into 3 splits for cross-validation. OATS labels 59 traffic atomic activity classes. We follow the same experimental protocol to train and evaluate the 35 classes. The dimensions of an image are  $1200 \times 1920$  pixels. We downsample the image size to  $224 \times 224$ , as in [3].

**TACO.** We construct the TACO dataset in the CARLA simulator [18] to attain a balanced distribution encompassing 64 classes of traffic atomic activities. We leverage three approaches to collect data: auto-pilot, scenario-runner [2], and automatic scenario generation [37]. The TACO dataset comprises 5178 clips. Among them, 1148 are used for testing. We collect image and instance segmentation data with a size of  $512 \times 1536$  pixels, and subsequently downsample each frame to  $256 \times 768$  as input to the models.

**nuScenes [9].** We annotate the *train\_val* set of nuScenes dataset with atomic activities. The *train\_val* set comprises 850 videos. Our annotation protocol yields a total of 426 short clips, each comprising 16 frames. The annotated nuScenes encompasses 42 classes of atomic activity. Please refer to the supplementary material for additional details on the dataset construction and annotation.

### 4.2. Baselines

**Video-level models.** We implement our baselines using the PyTorchVideo [21] library to construct different video-level architectures, including I3D [12], X3D [23], CSN [57], SlowFast [24], MViT [22], and VideoMAE [56]. Note that we modify the models to perform multi-label prediction with one linear layer that outputs  $N_{c1}$  class channels.

**Object-aware models.** The models [3, 6, 60] first extract object features via the RoI-Align [28] and learn the relations between objects. ORN [6], ARG [60], and OATS [3] uses MLP, Graph Convolutional Network (GCN) [35], and spatical-temporal GCN [44] to enhance object features. OATS also uses an additional graph branch to encode low-level object trajectories. We use Deep OC-SORT [42] to extract trajectories of objects. We follow OATS by setting the

number of tracklets as 20 for the input of the models. To encourage the object-aware models [6, 60] to learn the context information, we concatenate object features with the global features, which are obtained by convolving X3D [23] features with a 3D kernel size of 1. These models then output multi-class predictions for each proposal by a linear layer with a channel number of  $N_{cl}$  plus one negative class. Note that we do not implement the model proposed in OATS [3] due to the absence of source code.

**Slot-based models.** The models take features from the last ConvBlock of the encoders [23, 27] as input. Note that we remove the last projection layer of X3D, which makes the size of some slot-based models slightly smaller than X3D. We re-implement existing slot-attention-based methods for videos [5, 20, 34, 65] to use the proposed “allocated slots” for multi-label atomic activity recognition. Specifically, SAVi [34] randomly initialize slots for each forward pass from a learned distribution, while MO [5] and Slot-VPS [65] treat slots as learnable queries [11]. For the slot update mechanism, SAVi and MO recurrently update slots in time, while Slot-VPS uses an extra self-attention [58] to update slots across frames. We then use their slots in the last frame as the input to the classifier.

**Table 1.** Quantitative results on the OATS dataset. “Seq” denotes the input sequence length. The symbol ‡ denotes the re-implementation of slot-based methods, where each slot is allocated to a specific atomic activity class.  $S1$ ,  $S2$ , and  $S3$  represents the three defined splits in the OATS dataset.

Method	Backbone	Seq	Pretrain	S1	S2	S3	mAP
CSN [3, 57]	ResNet152	32	IG65M	12.1	12.6	12.9	12.5
TPN [3, 63]	ResNet50	32	ImageNet	11.6	13.3	12.9	12.6
SlowOnly [3, 24]	ResNet50	32	ImageNet	11.2	14.7	12.9	13.0
SlowFast [3, 24]	ResNet50	32	None	10.8	15.1	14.5	13.5
I3D (NL) [3, 59]	ResNet50	32	ImageNet	11.9	15.5	14.0	13.8
I3D [3, 12]	ResNet50	32	ImageNet	11.8	14.3	16.8	14.3
ORN [3, 6]	ResNet50	32	ImageNet	16.8	13.4	18.1	16.1
ARG [3, 60]	Inceptionv3	32	ImageNet	20.2	21.3	19.3	20.3
OATS [3]	Inceptionv3	32	ImageNet	24.3	28.6	27.2	26.7
SAVi‡ [20, 34]	ResNet50	32	ImageNet	21.1	22.4	22.5	22.0
MO‡ [5]	ResNet50	32	ImageNet	14.3	15.6	18.2	16.0
Slot-VPS‡ [65]	ResNet50	32	ImageNet	15.7	17.8	17.2	16.9
Action-slot (Ours)	ResNet50	32	ImageNet	<b>26.6</b>	<b>28.6</b>	<b>30.8</b>	<b>28.6</b>
I3D [12]	ResNet50	8	Kinetics-400	21.7	24.6	24.4	23.6
X3D [23]	N/A	16	Kinetics-400	30.4	33.2	30.6	31.4
CSN [57]	ResNet101	32	Kinetics-400	43.1	47.1	44.3	44.8
SlowFast [24]	ResNet50	16	Kinetics-400	36.1	36.6	34.2	35.6
ORN [6]	X3D	16	Kinetics-400	19.3	22.5	23.6	21.8
ARG [60]	X3D	16	Kinetics-400	24.8	25.9	29.3	26.7
SAVi‡ [20, 34]	X3D	16	Kinetics-400	19.0	22.1	21.6	20.9
MO‡ [5]	X3D	16	Kinetics-400	25.3	25.0	24.2	24.8
Slot-VPS‡ [65]	X3D	16	Kinetics-400	24.7	24.6	25.5	24.9
Action-slot (Ours)	X3D	16	Kinetics-400	<b>48.1</b>	<b>47.7</b>	<b>48.8</b>	<b>48.2</b>

**Pretraining and backbone choice.** We first follow the setting of OATS and implement Action-slot and slot-based baselines with ResNet50 [27] pretrained on ImageNet [17]. To further study the effect of pre-trained models, we train models with state-of-the-art 3D backbones [12, 23, 24], pre-

**Table 2.** Quantitative results on the TACO dataset. “Seq” denotes the input sequence length.  $C$ ,  $K$ ,  $P$ ,  $C+$ ,  $K+$ , and  $P+$  denote activities involved with different types of road users. The symbol ‡ denotes the re-implementation of slot-based methods, where each slot is allocated to a specific atomic activity class.

Method	Para. (M)	Seq	C	K	P	C+	K+	P+	mAP
I3D [12]	27.3	8	27.3	19.4	30.5	34.6	33.6	34.8	29.7
X3D [23]	3.0	16	37.5	20.3	34.6	56.3	51.5	38.8	38.3
CSN [57]	21.4	32	43.5	35.5	43.0	52.5	46.1	43.4	44.0
SlowFast [24]	33.7	16	33.9	19.7	36.7	39.9	42.6	41.0	35.2
MViT [22]	36.6	16	21.4	13.8	26.3	43.7	30.0	33.8	27.9
VideoMAE [22]	57.9	16	30.6	18.6	27.1	51.6	33.0	37.4	33.1
ORN [6]	4.8	16	25.5	15.8	24.6	31.6	19.8	13.9	22.2
ARG [60]	12.2	16	27.8	15.0	27.2	35.6	13.7	20.0	23.2
SAVi‡ [20, 34]	2.3	16	19.2	13.6	27.5	23.3	25.9	35.9	23.3
MO‡ [5]	2.3	16	34.2	24.9	39.2	38.3	39.4	37.7	35.3
Slot-VPS‡ [65]	3.5	16	31.9	21.3	32.0	51.9	44.7	31.3	36.0
Action-slot (Ours)	<b>2.3</b>	<b>16</b>	<b>48.1</b>	<b>41.2</b>	<b>49.2</b>	<b>70.1</b>	<b>62.6</b>	<b>52.8</b>	<b>54.4</b>

trained on Kinetics-400 [33], for the experiments on both OATS and TACO datasets. Specifically, we use X3D [23] as the backbone by default because of the compact model size. Please refer to the supplementary material for the additional implementation details.

### 4.3. Implementation Details

**Background masks.** We extract background masks from the instance segmentation collected in TACO. For OATS and nuScenes, we use an off-the-shelf DeepLabV3+ [14].

**Metrics.** We follow the common practice in multi-label action recognition [3, 26, 52, 64] to report mean average precision (mAP).

**Visualization.** We visualize attention maps of each atomic activity according to the allocated action slots to validate if Action-slot learns action-centric representations. We keep pixels with an attention score greater than 0.5 on OATS and 0.2 on TACO.

### 4.4. Main Results

**OATS.** In Table 1, we show experimental comparisons with video-level-based, object-aware, and slot-attention-based methods. In the top group, we use the ImageNet pre-trained ResNet50, following the evaluation protocol designed in OATS [3]. We further compare all methods with advanced visual backbones pretrained from Kinetics-400 [33] and report the corresponding results in the bottom group. In both ImageNet and Kinetics pretrained settings, Action-slot achieves state-of-the-art performance on the OATS dataset. Specifically, Action-slot improves the performance of atomic activity recognition by identifying the region of interest via the proposed slot attention mechanism, without explicit perception guidance.

**TACO.** We establish a new benchmark on the proposed TACO dataset to validate the effectiveness of Action-slot for a wide range of atomic activities, which is challeng-

**Table 3.** Ablation study of Action-slot. The reported results are the average of the 3 splits in the OATS dataset.

ID	Allocated	Update	BG Slot	$L_{bg}$	$L_{neg}$	mAP
1		parallel				10.8
2	✓	recurrent				24.8
3	✓	parallel				42.7
4	✓	parallel	✓			40.8
5	✓	parallel	✓	✓		43.6
6	✓	parallel	✓		✓	43.0
7	✓	parallel	✓	✓	✓	<b>48.2</b>

ing for the OATS dataset [3] due to its imbalanced distribution. In the top group of Table 2, we evaluate two recent transformer-based models, i.e., MViT [22] and VideoMAE [56]. The models show worse performance compared to other video-level methods, which could be attributed to their data-hungry nature [19].

Object-aware models [3, 6, 60] show suboptimal performance, particularly for atomic activities related to grouped road users, as reported in the middle group. The findings indicate that the relation modeling of objects learned from the MLP and GCN [35] is not effective for the task because it requires a holistic understanding of contextual information and motions of objects.

In the bottom group of Table 2, we re-implement the object-centric slot-based models with allocated slots (denoted by ‡). Action-slot outperforms all models by a large margin on activities with all types of road users because of our parallel architecture design and the proposed attention guidance. Particularly, Action-slot excels in activities involved with grouped road users. It is worth noting that the Action-slot is powerful yet efficient in terms of model size. The results validate the effectiveness of action-centric representations over the existing action recognition models and object-centric representations.

#### 4.5. Ablation Study

We conduct ablation studies in Table 3 to understand the effect of our designed components.

**Non-Allocated vs. Allocated.** We show that adding allocated slots significantly improves the performance of **Non-Allocated** by 31.9% (ID 1 vs. 3). The results suggest the need to revisit the design of slot attention for our task.

**Recurrent vs. Parallel.** We demonstrate that updating action slots in a parallel fashion obtains a 17.9% performance gain (ID 2 vs. 3), compared to conventional methods using a recurrent manner.

**Background slot and attention guidance  $L_{bg}$ .** We find that simply adding a non-allocated background slot without mask supervision slightly harms the performance (ID 3 vs. 4), which is different from the observation in previous slot attention literature [5, 65]. Our insight is that the non-allocated slot may distract the action slots’ attention. With the background mask supervision  $L_{bg}$ , the model shows ad-

**Table 4.** Comparisons of Action-slot and object-level guidance on TACO across different numbers of road users (denoted as  $N$ ) present in a video. BG and Neg denote background slot and regularization in our method.

	$N \leq 5$	$5 < N \leq 15$	$N > 15$
object	49.4	47.7	43.5
BG+Neg	55.2	50.9	46.3

**Table 5.** Results of various backbones for Action-slot on TACO.

	mAP
I3D	29.7
Action-slot	37.6 (+7.9)
X3D [23]	37.8
Action-slot	<b>54.4 (+16.6)</b>
SlowFast [24]	35.2
Action-slot	46.7 (+11.5)

ditional performance gain compared to the one only using the allocated slot (ID 3 vs. 5).

**Action slots regularization  $L_{neg}$ .** We regularize action slots with the  $L_{neg}$  term to discourage action slots allocated to negative classes from attending any regions in spatial-temporal features. This design enhances the likelihood of other action slots allocated to positive classes identifying the region of interest more effectively. We demonstrate that the inclusion of the regularization term  $L_{neg}$  further improves the one only using the allocated slot (ID 3 vs. 6).


The final model (ID 7 in Table 3) that incorporates allocated slots, parallel updating, background slot with attention guidance, and action slot regularization, performs the best on OATS, which validates the efficacy of Action-slot.

## 5. More Analysis and Discussions

**Action-slot vs. Object-level guidance.** We conduct experiments on the TACO dataset to evaluate the impact of different numbers of road users in videos with respect to different attention guidance mechanisms in our Action-slot framework. Specifically, we compare the two attention guidance mechanisms discussed in Section 3.3 and the object-level guidance. For object-level guidance, we use instance segmentation collected in TACO to obtain masks for vehicles and pedestrians. Hungarian matcher is used to assign object masks to slots, following MO [5]. Table 4 shows that Action-slot with strong ground truth object-level supervisions performs worse, compared to the proposed attention guidance. The observation indicates that only considering object cues may not effectively tackle our task because not every object is always involved in the same atomic activity.

**Action-slot with different backbones.** Table 5 reports the performance of Action-slot using different backbone. We choose I3D [12], X3D [23], and SlowFast [24] for experiments. Action-slot significantly improves all 3D CNN methods, showing great generalization ability. Among them, X3D achieves the most improvement. We hypothesize that this is because X3D’s features retain the original temporal dimension, affording Action-slot greater spatial-temporal information to analyze atomic activities. In contrast, other backbone encoders downsample the temporal dimension; for instance, SlowFast takes a video sequence of





**Figure 5.** Visualization of attention maps learned from OATS. Colored masks represent the attention of the activity slots **Z4-Z3:C** and **C2-C1:P+**. Note that, while MO successfully predicts the activity **C2-C1:P+**, the corresponding attention scores are very low.



**Figure 6.** Visualization of attention maps learned from TACO. Colored masks are from the allocated slots of activity **Z4-Z1:K** and **Z1-Z2:K+**. Note that **Z4-Z1:K** appears twice in the 1st and 5th frames independently.

**Table 6.** Effectiveness of using our TACO dataset for representation pretraining and then finetuning on real-world OATS and nuScenes datasets. Kinetics denotes finetuning from the Kinetics-400 pretrained weights.

	OATS		nuScenes	
	Kinetics	+TACO	Kinetics	+TACO
X3D [23]	31.4	34.9 (+3.5)	19.8	27.8 (+8.0)
ARG [60]	26.7	31.3 (+4.6)	12.2	17.0 (+4.8)
Action-slot	48.2	<b>59.5</b> (+11.3)	23.6	<b>32.3</b> (+8.7)

length 16 as input and reduces the temporal dimension to 4.

**Representation pretraining on TACO.** We demonstrate that the proposed TACO dataset can enhance the efficacy of atomic activity recognition on real-world datasets through representation pretraining. Specifically, we first pretrain the model on TACO and then fine-tune it on real-world datasets, including OATS [3] and a newly annotated nuScenes dataset [9]. We report the average mAP of the three splits in the OATS dataset in Table 6. Due to the large-scale and balanced distribution of TACO, all models show a significant improvement on both real-world datasets when pretrained on TACO. Notably, Action-slot outperforms all baselines, showcasing its robustness and generalizability.

**Qualitative Results.** We visualize the attention maps learned using the re-implemented MO [5] and Action-slot on the OATS and TACO datasets in Figure 5 and Figure 6, respectively. Action-slot demonstrates the capability to decompose multiple atomic activities from intricate scenes in both datasets. In Figure 5, Action-slot localizes the group activity **C2-C1:P+**, while MO [5] fails to attend any relevant regions. Moreover, Action-slot recognizes two independent **Z4-Z1:K** activities, instead of predicting them as a

group activity, i.e., **Z4-Z1:K+** in Figure 6. The evidence justifies the capability of Action-slot that learns action-centric representations. More interestingly, Action-slot attends to relevant regions when road users initiate an action and cease attention once the action is completed, as seen in the case of the bicyclists in the fourth frame of Figure 6.

## 6. Conclusions

We present Action-slot, a slot attention-based approach that learns visual action-centric representations, capturing both motion and contextual information. We introduce several crucial designs in slot attentions that enable decomposing multiple atomic activities in videos, without the need for explicit perception guidance. We conduct comprehensive experiments to validate the proposed method on both OATS and our TACO dataset, comparing it against various action recognition baseline methods. In addition, we demonstrate that Action-slot learns meaningful attention maps, identifying objects in actions without explicit object-level guidance. Lastly, we show that the TACO dataset enhances the efficacy of multi-label atomic activity recognition on real-world datasets through representation pretraining.

**Limitation.** Although Action-slot demonstrates favorable quantitative and qualitative performance, we observe that it is still challenging to handle cases when two activities occlude with each other. In these cases, the corresponding action slots may confuse where they should attend. This observation is closely relevant to the tracking task under occlusion, where frequent ID switches are often observed. We hope our findings can encourage the community to explore more advanced designs for action-centric representations.



## References

- [1] CARLA Autonomous Driving Challenge. <https://carlachallenge.org/>, 2022. **1**
- [2] ScenarioRunner for CARLA. [https://github.com/carla-simulator/scenario\\_runner](https://github.com/carla-simulator/scenario_runner), 2023. **5**
- [3] Nakul Agarwal and Yi-Ting Chen. Ordered atomic activity for fine-grained interactive traffic scenario understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8624–8636, 2023. **1, 2, 3, 5, 6, 7, 8**
- [4] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A Video Vision Transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6836–6846, 2021. **3, 4**
- [5] Zhipeng Bao, Pavel Tokmakov, Allan Jabri, Yu-Xiong Wang, Adrien Gaidon, and Martial Hebert. Discovering Objects that Can Move. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11789–11798, 2022. **2, 3, 4, 5, 6, 7, 8**
- [6] Fabien Baradel, Natalia Neverova, Christian Wolf, Julien Mille, and Greg Mori. Object Level Visual Reasoning in Videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. **2, 5, 6, 7**
- [7] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *International Conference on Machine Learning (ICML)*, 2021. **3, 4**
- [8] Ondrej Biza, Sjoerd van Steenkiste, Mehdi SM Sajjadi, Gamaleldin F Elsayed, Aravindh Mahendran, and Thomas Kipf. Invariant Slot Attention: Object Discovery with Slot-Centric Reference Frames. *International Conference on Machine Learning (ICML)*, 2023. **3**
- [9] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A Multi-modal Dataset for Autonomous Driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. **2, 5, 8**
- [10] Yulong Cao, Chaowei Xiao, Anima Anandkumar, Danfei Xu, and Marco Pavone. AdvDO: Realistic Adversarial Attacks for Trajectory Prediction. In *European Conference on Computer Vision (ECCV)*. Springer, 2022. **1, 3**
- [11] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end Object Detection with Transformers. In *European Conference on Computer Vision (ECCV)*, pages 213–229. Springer, 2020. **6**
- [12] 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 (CVPR)*, pages 6299–6308, 2017. **2, 3, 4, 5, 6, 7**
- [13] Chao-Yeh Chen, Wongun Choi, and Manmohan Chandraker. Atomic Scenes for Scalable Traffic Scene Recognition in Monocular Videos. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9, 2016. **1, 3**
- [14] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. **6**
- [15] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014. **4**
- [16] Wongun Choi, Khuram Shahid, and Silvio Savarese. What are they doing?: Collective activity classification using spatio-temporal relationship among people. In *IEEE International Conference on Computer Vision Workshops (ICCV-W)*, pages 1282–1289, 2009. **3**
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A Large-scale Hierarchical Image Database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. **6**
- [18] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16, 2017. **2, 3, 5**
- [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. In *International Conference on Learning Representations (ICLR)*, 2021. **7**
- [20] Gamaleldin Elsayed, Aravindh Mahendran, Sjoerd van Steenkiste, Klaus Greff, Michael C Mozer, and Thomas Kipf. SAVi++: Towards End-to-End Object-centric Learning from Real-world Videos. *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. **2, 3, 4, 5, 6**
- [21] Haoqi Fan, Tullie Murrell, Heng Wang, Kalyan Vasudev Alwala, Yanghao Li, Yilei Li, Bo Xiong, Nikhila Ravi, Meng Li, Haichuan Yang, Jitendra Malik, Ross Girshick, Matt Feiszli, Aaron Adcock, Wan-Yen Lo, and Christoph Feichtenhofer. PyTorchVideo: A Deep Learning Library for Video Understanding. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. <https://pytorchvideo.org/>. **5**
- [22] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6824–6835, 2021. **2, 3, 4, 5, 6, 7**
- [23] Christoph Feichtenhofer. X3D: Expanding Architectures for Efficient Video Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 203–213, 2020. **2, 3, 4, 5, 6, 7, 8**
- [24] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast Networks for Video Recognition. In

- Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6202–6211, 2019. 2, 3, 4, 5, 6, 7
- [25] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 3
- [26] 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 (CVPR)*, pages 6047–6056, 2018. 3, 6
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 4, 6
- [28] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2961–2969, 2017. 2, 5
- [29] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1971–1980, 2016. 3
- [30] Ashesh Jain, Hema S. Koppula, Bharad Raghavan, Shane Soh, and Ashutosh Saxena. Car that Knows Before You Do: Anticipating Maneuvers via Learning Temporal Driving Models. In *IEEE International Conference on Computer Vision (ICCV)*, 2015. 1
- [31] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://csrcv.ucf.edu/THUMOS14/>, 2014. 3
- [32] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2901–2910, 2017. 3
- [33] 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. 3, 6
- [34] Thomas Kipf, Gamaleldin F. Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonckhowski, Alexey Dosovitskiy, and Klaus Greff. Conditional Object-Centric Learning from Video. In *International Conference on Learning Representations (ICLR)*, 2022. 2, 3, 4, 5, 6
- [35] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 5, 7
- [36] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: A Large Video Database for Human Motion Recognition. In *2011 International Conference on Computer Vision (ICCV)*, pages 2556–2563, 2011. 3
- [37] Chi-Hsi Kung, Chieh-Chi Yang, Pang-Yuan Pao, Shu-Wei Lu, Pin-Lun Chen, Hsin-Cheng Lu, and Yi-Ting Chen. Riskbench: A scenario-based benchmark for risk identification. *arXiv preprint arXiv:2312.01659*, 2023. 5
- [38] Max Guangyu Li, Bo Jiang, Zhengping Che, Xuefeng Shi, Mengyao Liu, Yiping Meng, Jieping Ye, and Yan Liu. DBUS: Human Driving Behavior Understanding System. In *ICCV Workshops*, pages 2436–2444, 2019. 3
- [39] Dahua Lin, Sanja Fidler, Chen Kong, and Raquel Urtasun. Visual Semantic Search: Retrieving Videos via Complex Textual Queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2014. 1, 3
- [40] Bingbin Liu, Ehsan Adeli, Zhangjie Cao, Kuan-Hui Lee, Abhijeet Shenoi, Adrien Gaidon, and Juan Carlos Nieves. Spatiotemporal relationship reasoning for pedestrian intent prediction. pages 3485–3492, 2020. 1, 3
- [41] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric Learning with Slot Attention. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:11525–11538, 2020. 2, 3, 4
- [42] Gerard Maggolino, Adnan Ahmad, Jinkun Cao, and Kris Kitani. Deep OC-SORT: Multi-Pedestrian Tracking by Adaptive Re-Identification. *arXiv preprint arXiv:2302.11813*, 2023. 5
- [43] Srikanth Malla, Behzad Dariush, and Chiho Choi. TITAN: Future Forecast using Action Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11186–11196, 2020. 1, 3
- [44] Abdullh Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. Social-stgcn: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14424–14432, 2020. 5
- [45] Milind Naphade, Shuo Wang, David C. Anastasiu, Zheng Tang, Ming-Ching Chang, Yue Yao, Liang Zheng, Mohammed Shaiqur Rahman, Meenakshi S. Arya, Anuj Sharma, Qi Feng, Vitaly Ablavsky, Stan Sclaroff, Pranamesh Chakraborty, Sanjita Prajapati, Alice Li, Shangru Li, Krishna Kunadharaju, Shenxin Jiang, and Rama Chellappa. The 7th AI City Challenge. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2023. 1, 3
- [46] Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, and Kate Saenko. Toward Driving Scene Understanding: A Dataset for Learning Driver Behavior and Causal Reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7699–7707, 2018. 1, 3
- [47] Amir Rasouli, Iuliia Kotseruba, Toni Kunic, and John K. Tsotsos. PIE: A Large-Scale Dataset and Models for Pedestrian Intention Estimation and Trajectory Prediction. In *The*

- IEEE International Conference on Computer Vision (ICCV)*, 2019. 1
- [48] Davis Rempe, Jonah Philion, Leonidas J. Guibas, Sanja Fidler, and Or Litany. Generating useful accident-prone driving scenarios via a learned traffic prior. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3
- [49] Davis Rempe, Zhengyi Luo, Xue Bin Peng, Ye Yuan, Kris Kitani, Karsten Kreis, Sanja Fidler, and Or Litany. Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 3
- [50] Mehdi S. M. Sajjadi, Daniel Duckworth, Aravindh Mahendran, Sjoerd van Steenkiste, Filip Pavetic, Mario Lucic, Leonidas J Guibas, Klaus Greff, and Thomas Kipf. Object scene representation transformer. In *Advances in Neural Information Processing Systems*, pages 9512–9524. Curran Associates, Inc., 2022. 3
- [51] Sean Segal, Eric Kee, Wenjie Luo, Abbas Sadat, Ersin Yumer, and Raquel Urtasun. Universal Embeddings for Spatio-Temporal Tagging of Self-Driving Logs. In *Conference on Robot Learning (CoRL)*, 2020. 1, 3
- [52] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 510–526. Springer, 2016. 3, 6
- [53] Gurkirt Singh, Stephen Akrigg, Manuele Di Maio, Valentina Fontana, Reza Javanmard Alitappeh, Salman Khan, Suman Saha, Kossar Jeddisaravi, Farzad Yousefi, Jacob Culley, et al. ROAD: The ROad event Awareness Dataset for Autonomous Driving. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1036–1054, 2022. 1, 3
- [54] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 3
- [55] Ahmed Taha, Yi-Ting Chen, Teruhisa Misu, Abhinav Shrivastava, and Larry Davis. Boosting Standard Classification Architectures Through a Ranking Regularizer. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2020. 1, 3
- [56] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural Information Processing Systems*, 2022. 2, 3, 4, 5, 7
- [57] Du Tran, Heng Wang, Lorenzo Torresani, and Matt Feiszli. Video Classification with Channel-separated Convolutional Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5552–5561, 2019. 2, 3, 5, 6
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Conference on Neural Information Processing Systems*, 30, 2017. 4, 6
- [59] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. 6
- [60] Jianchao Wu, Limin Wang, Li Wang, Jie Guo, and Gangshan Wu. Learning Actor Relation Graphs for Group Activity Recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2, 5, 6, 7, 8
- [61] Zihao Xiao, Allan Yuille, and Yi-Ting Chen. Learning Road Scene-level Representations via Semantic Region Prediction. In *Conference on Robot Learning (CoRL)*, 2022. 1, 3
- [62] Chejian Xu, Wenhao Ding, Weijie Lyu, Zuxin Liu, Shuai Wang, Yihan He, Hanjiang Hu, Ding Zhao, and Bo Li. Safebench: A benchmarking platform for safety evaluation of autonomous vehicles. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 1, 3
- [63] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 591–600, 2020. 6
- [64] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 2017. 3, 6
- [65] Yi Zhou, Hui Zhang, Hana Lee, Shuyang Sun, Pingjun Li, Yangguang Zhu, ByungIn Yoo, Xiaojuan Qi, and JaeJoon Han. Slot-VPS: Object-centric Representation Learning for Video Panoptic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3093–3103, 2022. 3, 4, 6, 7