

## How can objects help action recognition?

Xingyi Zhou Anurag Arnab Chen Sun Cordelia Schmid  
 Google Research

{zhouxy, aarnab, chensun, cordelias}@google.com

### Abstract

Current state-of-the-art video models process a video clip as a long sequence of spatio-temporal tokens. However, they do not explicitly model objects, their interactions across the video, and instead process all the tokens in the video. In this paper, we investigate how we can use knowledge of objects to design better video models, namely to process fewer tokens and to improve recognition accuracy. This is in contrast to prior works which either drop tokens at the cost of accuracy, or increase accuracy whilst also increasing the computation required. First, we propose an object-guided token sampling strategy that enables us to retain a small fraction of the input tokens with minimal impact on accuracy. And second, we propose an object-aware attention module that enriches our feature representation with object information and improves overall accuracy. Our resulting model, ObjectViViT, achieves better performance when using fewer tokens than strong baselines. In particular, we match our baseline with 30%, 40%, and 60% of the input tokens on *SomethingElse*, *Something-something v2*, and *Epic-Kitchens*, respectively. When we use ObjectViViT to process the same number of tokens as our baseline, we improve by 0.6 to 4.2 points on these datasets.

### 1. Introduction

Video understanding is a central task of computer vision and great progresses have been made recently with transformer-based models which interpret a video as a sequence of spatio-temporal tokens [1, 3, 13, 16, 30, 55]. However, videos contain a large amount of redundancy [53], especially when there is little motion or when backgrounds remain static. Processing videos with all these tokens is both inefficient and distracting. As objects conduct motions and actions [19, 47], they present an opportunity to form a more compact representation of a video, and inspire us to study how we can use them to understand videos more accurately and efficiently. Our intuition is also supported biologically that we humans perceive the world by concentrating our focus on key regions in the scene [18, 42].



Figure 1. **How can objects help action recognition?** Consider this action of picking up a bowl on a countertop packed with kitchenware. Objects provide information to: (1) associate image patches (colorful) from the same instance, and identify candidates for interactions; (2) selectively build contextual information from the redundant background patches (dark).

In this paper, we explore how we can use external object information in videos to improve recognition accuracy, and to reduce redundancy in the input (Figure 1). Current approaches in the literature have proposed object-based models which utilize external object detections to improve action recognition accuracy [21, 35]. However, they aim to build architectures to model objects and overall bring a notable computational overhead. We show besides gaining accuracy, objects are extremely useful to reduce token redundancy, and give a more compact video representation. Fewer tokens also enable stronger test strategies (e.g., multi-crop, longer videos), and overall improve accuracy further. On the other hand, prior work on adaptive transformers [36, 45, 48, 50] dynamically reduce the number of tokens processed by the network conditioned on the input. Since these methods learn token-dropping policies end-to-end without external hints, there can be a chicken-and-egg problem: we need good features to know which token to drop, and need many tokens to learn good features. As a result, they usually suffer from a performance decrease when dropping tokens. We show we can perform both goals, namely reducing tokens processed by the transformer, as well as improving accuracy, within a unified framework.

Concretely, we propose an **object-based video vision transformer**, ObjectViViT. As shown in Figure 2, we first propose an object-guided token sampling strategy (OGS) that uses object locations to identify foreground and background tokens. We retain relevant foreground tokens as is, and aggressively downsample background tokens be-




Figure 2. **Illustration of our object-based video vision transformer, ObjectViViT.** ObjectViViT takes raw video pixels and off-the-shelf object detections (bounding boxes) as input, and runs space-time attention on video tokens [1]. We use the detection boxes in two ways. (1): we use object locations to downsample the patch tokens before running transformers (Object-Guided Token Sampling, more details in Figure 3). (2): we run a customized attention module that creates object tokens from object-patch relations and uses them to enhance patch features (Object-Aware Attention Module, more details in Figure 4).

fore forwarding them to the transformer module. Secondly, to fully leverage the relation between objects and the unstructured spatial-temporal patches, we introduce an object-aware attention module (OAM). This attention module first creates object tokens by grouping patch tokens from the same object using an object-weighted pooling, and then applies space-time attention on the concatenated object and patch tokens. This way, patch features are augmented with their related object information. Both OGS and OAM are complementary. They can be used individually to improve either token-compactness or accuracy, or can be used together to get benefits of both.

We validate our method with extensive experiments on SomethingElse [32], Something-Something [17], and the Epic Kitchens datasets [8]. Using our object-guided token sampling, we find that we can process 60% ~ 90% of the input tokens without losing any accuracy. And by using our object-aware attention module alone, we outperform a competitive ViViT [1] baseline by 0.6 to 2.1 points. Combining both modules, ObjectViViT improves token-compactness and accuracy even further, matching baseline performance by processing 30%, 40%, and 60% of the input tokens for the three datasets, respectively. Finally, under the same number of processed tokens but a higher temporal resolution, our model with dropped tokens improve upon baselines by up to 4.2 points. Our code is released at <https://github.com/google-research/scenic>.

## 2. Related Work

**Object-based video representation** aims to integrate object information in videos [12, 31, 39] and is a popular and promising direction in video recognition. Wang et al. [47] extract RoI features from a 3D CNN feature maps using off-the-shelf detectors, and build a graph neural network on RoI features alone. LFB [51] and Object Transformer [52] load

off-the-shelf object features from detectors in video backbones to as feature banks for long video. Recently, with the advance of transformer-based architectures, ORViT [21] proposes to crop [20] object regions as a new object tokens and attach them to pixel tokens. ObjectLearner [57] fuses an object-layout stream and a pixels stream using an object-to-pixel transformer. Most of these existing methods focus on enhancing video features, and so as ours. However instead of representing objects as cropped boxes, we obtain object features using a simpler heatmap-weighted pool. In sharp contrast to these methods, we in addition use objects to decrease transformer token redundancy by using our object-guided token downsampling.

On a parallel direction, STIN [32] and STLT [35] propose object-layout only architectures without any visual inputs. These methods only work well with oracle object detections where motions can be perfectly encoded in objects. However both still report a significant performance drop compared to the pixel-based counterparts. Our work encodes object information inside a visual backbone by both our sampler and attention module, and retains full-pixel performance with very few visual tokens.

**Transformers for videos.** Transformers [11] are now the predominant architectures for video recognition. The vanilla vision transformer (ViT) [11] evenly divides images into non-overlapping tokens, and runs multi-head self-attention [44] over the tokens. TimeSFormer [3] and ViViT [1] extends ViT to videos, by introducing tube tokenization and efficient cross-time attention, i.e., axis-based space-time attention [3] or factorized attention [1]. MotionFormer [34] enhances space-time attention by an implicit trajectory attention module. MViT [13, 28] and VideoSwin Transformer [30] re-introduce resolution-pooling as in convolution nets in video transformers for efficiency.

Our model is built on top of vanilla space-time attention in ViViT [1] without factorization, which can directly ben-




Figure 3. **Details about our object-guided token sampling (OGS).** We first render the off-the-shelf object detections as center heatmaps [60]. After tokenization, we define the “object-scores” of each token as the sum of the heatmap values inside it. We define the tokens with top  $X\%$  scores as foreground, where  $X$  is a configurable parameter, and the rest as background. We keep all foreground tokens, and aggressively downsample the background tokens. This mechanism automatically prioritizes object coverage as all object centers have the highest score.

efit from Masked pretraining [43] and gives state-of-the-art results [15, 43]. We introduce objects as additional sources to improve both token-efficiency and accuracy.

**Token-efficiency in transformers.** Dropping tokens is a popular idea to make transformer models efficient [14, 26, 33, 36]. DynamicViT [36] learns a token scoring function end-to-end using Gumbel-Softmax [23]. A-ViT [56] adds a loss for the number of active tokens and encourages classifying images with fewer tokens. ToMe [5] merges tokens by measuring their similarity in test time. In videos, STTS [45] applies a Top-K operator first on time to select keyframes and then on manually designed spatial anchors to select key regions. AdaFocus [48–50] designs a region selection module using recurrent networks, with an early termination mechanism that stops processing after a certain confidence score.

All these methods learn the token scoring function in an unsupervised way without additional information source. As a result, they see a reduction in performance as information is reduced. For example, when STTS [45] reduces the tokens to 60%, accuracy degrades by 1.5%. In contrast, as we use objects to guide our token selection, we do not degrade, and even improve by 0.2% at this dropping ratio. However efficiency is a nuanced topic [9], and whilst our approach does not make the inference pipeline more efficient as it relies on additional object detectors, it does shine a light on how many tokens we eventually need to process with state-of-the-art transformer architectures. Furthermore, as we train our model while dropping tokens, it means that we can in fact improve overall training efficiency over standard video models which do not drop tokens.

### 3. Preliminary

Given a video  $V$  represented as a stack of pixels  $V \in \mathbb{R}^{T \times H \times W \times 3}$ , action classification aims to classify the video into one of the pre-defined action categories  $c \in \mathcal{C}$ . We build our framework on top of space-time ViViT [1], which is a straightforward extension of image ViT [11].

ViViT [1] first divides video pixels  $V$  into non-overlapping space-time tokens, where each token repre-

sents a pixel tube in size  $(dt, dh, dw)$ . This gives  $N = \lfloor \frac{T}{dt} \rfloor \times \lfloor \frac{H}{dh} \rfloor \times \lfloor \frac{W}{dw} \rfloor$  tokens. The RGB pixels inside each token are then mapped to a fixed dimension  $D$  by a linear layer, followed by adding positional embeddings. As a result, the input to the transformer is  $\mathbf{z}^0 = [z_1, z_2, \dots, z_N]$ , where  $z_i \in \mathbb{R}^D$  is a patch token. Each transformer layer updates the tokens  $\mathbf{z}$  with multi-head attention (MHA) [44] and an MLP:

$$\mathbf{y}^l = \text{MHA}(\mathbf{z}^l, \mathbf{z}^l, \mathbf{z}^l) + \mathbf{z}^l \quad (1)$$

$$\mathbf{z}^{l+1} = \text{MLP}(\mathbf{y}^l) + \mathbf{y}^l \quad (2)$$

Where  $\text{MHA}(q, k, v)$  computes attentions between queries  $q$ , keys  $k$ , and values  $v$ . MLP contains 2 linear layers. We omit layer norm [2] for simplicity. The output classification logits  $\mathbf{c} \in \mathbb{R}^{|\mathcal{C}|}$  is obtained after a linear layer on global average-pooled features of all tokens in the last layer.

## 4. Method

Besides video pixels  $V$ , our framework takes an additional input: object detections  $\mathcal{B} = \{ \{ (b_{t,o}, \delta_{t,o}) \}_{o=1}^{n_t} \}_{t=1}^T$ , where  $b_{t,o} = (x_{t,o}, y_{t,o}, w_{t,o}, h_{t,o})$  is the bounding box of the  $o$ -th object in frame  $t$ ,  $\delta_{t,o}$  is an optional object identity (when tracking is enabled), and  $n_t$  is the number of objects in frame  $t$ . The bounding boxes are usually from an off-the-shelf object detector [37], or can be oracle boxes from annotation for analytic benchmarks [32]. Next, we introduce ObjectViViT, which includes two components, Object-guided token sampling (OGS) and Object-aware attention module (OAM), to use these external detections for token-efficient and accurate action recognition.

### 4.1. Object-guided token sampling

We first use objects  $\mathcal{B}$  to downsample input tokens  $\mathbf{z}^0$ . Figure 3 gives an overview of this process. Our main insight is that tokens which are inside objects carry the motions in the scene and should be retained, and tokens in the background are mostly static and can be partially dropped due to their redundancy. We also aim to configure the dropping ratio to be able to control the token-accuracy trade-off.

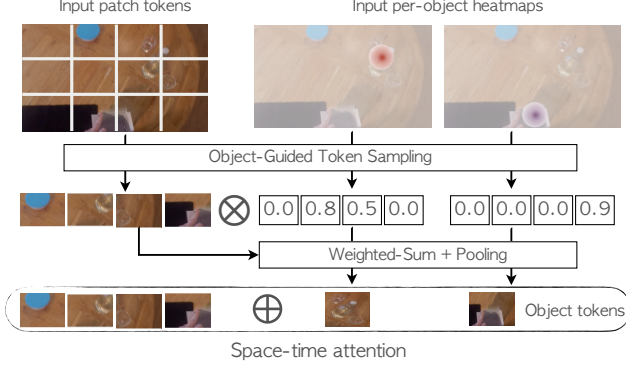


Figure 4. **Details about our object-aware attention module (OAM).** The inputs to the module are (downsampled) space-time patch tokens together with their object-scores with respect to each object instance. We only show one frame for clarity. For each object instance at each frame, we create a new object token by weighted pool between the pixel features and the object weights ( $\otimes$ ). We concatenate ( $\oplus$ ) the object tokens and patch tokens and conduct space-time attention over all tokens following ORViT [21]. The module outputs updated features of patch tokens, which are enhanced by object information.

To do so, we define a continuous token objectness score by measuring how close each token is to object centers. Specifically, we follow CenterNet [60] to render objects in each frame into a single class-agnostic center heatmap  $\mathbf{H}_t \in \mathbb{R}^{H \times W}$  in the original image resolution, where

$$\mathbf{H}_{t,i,j} = \max_{o \leq n^t} \exp\left(-\frac{(i-x_{t,o})^2 + (j-y_{t,o})^2}{2\sigma_{t,o}^2}\right), \quad (3)$$

and  $\sigma_{t,o}$  is a monotonic function of the object size that controls the Gaussian radius. We use the same tube size  $(dt, dh, dw)$  to project the heatmap into tubelets  $\tilde{\mathbf{H}} \in \mathbb{R}^N$ .

We simply select the top  $X\%$  of the tokens according to object-ness scores  $\tilde{\mathbf{H}}$ :

$$\mathbf{z}_{\text{fg}}^0 = \{z_n | \tilde{\mathbf{H}}_n \geq \tau\}, \quad \mathbf{z}_{\text{bg}}^0 = \mathbf{z}^0 - \mathbf{z}_{\text{fg}}^0 \quad (4)$$

where  $\tau$  is the  $X\%$ -th value of  $\tilde{\mathbf{H}}$ , and  $X$  is a configurable parameter to control the number of tokens.

With this foreground/ background token definition, we keep all the foreground tokens as is, and uniformly sample  $Y\% \times N$  tokens from the background. This gives the downsampled background tokens  $\hat{\mathbf{z}}_{\text{bg}}^0$ , where  $|\hat{\mathbf{z}}_{\text{bg}}^0| = Y\% \times N$ . Our final inputs are a concatenation of both tokens:

$$\hat{\mathbf{z}}^0 = [\mathbf{z}_{\text{fg}}^0, \hat{\mathbf{z}}_{\text{bg}}^0] \quad (5)$$

We therefore have  $|\hat{\mathbf{z}}^0| = (X + Y)\% \times N$  remaining tokens.

## 4.2. Object-aware attention

Next, we use objects  $\mathcal{B}$ , in transformer blocks to improve features (Figure 4). We again represent objects in

a frame as a center-heatmap [60], but here we use instance-specific heatmaps. I.e., for each object  $o$  in frame  $t$ , we render a heatmap  $\mathbf{H}_{o,t} \in \mathbb{R}^{H \times W}$  with a single Gaussian peak:  $\mathbf{H}_{o,t,i,j} = \exp\left(-\frac{(i-x_{t,o})^2 + (j-y_{t,o})^2}{2\sigma_{t,o}^2}\right)$ . Each  $\mathbf{H}_{o,t}$  highlights the pixel regions of object  $o$  in frame  $t$ . After tokenization and the token-sampling following Equation 5, the downsampled per-instance heatmap  $\hat{\mathbf{H}}_{o,t}$  represents the affinity score between object  $o$  and each remaining token in frame  $t$ . These affinity scores naturally serve as weighting scores to aggregate object features, and we directly use them to create object tokens  $\mathbf{w}_{o,t}$  for each object at each frame.

$$\mathbf{w}_{o,t}^l = \text{MaxPool}(\text{MLP}(\hat{\mathbf{H}}_{o,t} * \hat{\mathbf{z}}_t^l)), \quad (6)$$

where  $l$  is the transformer layer index,  $\hat{\mathbf{z}}_t^l$  is the (downsampled) patch tokens of frame  $t$  at layer  $l$ , and MLP are two linear layers following ORViT [21]. Optionally, when object identity information is available (through tracking), we encode object identity via a learnable identity embedding  $\mathbf{e}_o \in \mathbb{R}^D$  shared across time:  $\mathbf{w}_{o,t}^l = \mathbf{w}_{o,t}^l + \mathbf{e}_o^l$ .

We next concatenate the object tokens and the (downsampled) patch tokens and keys and values, and use them to update the query patch tokens in Equation 1:

$$\mathbf{y}^l = \text{MHA}(\mathbf{z}^l, [\mathbf{z}^l, \mathbf{w}^l], [\mathbf{z}^l, \mathbf{w}^l]) + \mathbf{z}^l \quad (7)$$

Following ORViT [21], we observe inserting the object tokens at 3 layers (the 2, 6, 11-th layer) of the transformer is sufficient. All other layers use the vanilla self-attention in Equation 1.

**Comparison to ORViT [21].** Our object-aware attention module is inspired by the Object-Region Attention of ORViT [21]. We both create object tokens from pixels and involve them in cross-attention computation. ORViT uses ROIAlign [20] to aggregate object features, which requires the input to be a full rectangular grid. They also use an axis-aligned trajectory attention backbone [34] which requires the same shape across frames. We use a simpler token aggregating function (Equation 5) that does not have these constraints and works for any input shape. We also don't use the bounding box layout module in ORViT [21]. We will show our formulation is simpler, more flexible, and performs competitively.

## 5. Experiments

We evaluate ObjectViViT for action classification on Something-something v2 [17], SomethingElse [32], and Epic Kitchen [8] datasets.

**Something-something v2 (SSv2)** [17] is a large-scale action classification dataset consisting of  $\sim 169\text{K}$  training videos and  $\sim 25\text{K}$  validation videos. The videos are 2-4 seconds long, focusing on hand-object interaction actions,




Figure 5. **Main results.** We show action classification results with our proposed **Object-aware attention** module and **Object-guided token sampling** under different numbers of sampled tokens ( $(X + Y)\%$ , the ratio with respect to all tokens) on SomethingElse (a), SSv2 (b), and Epic Kitchens (c) dataset. We show results with inferred boxes from detectors (**Ours**) or with **oracle boxes**. Our baseline uses Space-time ViViT [1] with VideoMAE pretraining [43]. The models are evaluated under single-crop and single-view testing by default. Using all tokens (x-axis at 1.0), our Object-aware attention improves the baseline by 0.6 to 2.1 points. With fewer tokens, our model matches the baseline performance with 30%, 40%, and 60% of tokens on the three datasets, respectively. We also show an optimal model (★) with 50% tokens and 2-view testing, which achieves better accuracy than the full model under the same number of processed tokens. With ground truth boxes, we only need 10% tokens to match the baseline’s performance.

with in total 174 action classes. The action classes are agnostic to object classes. **Something-Else** [32] is a re-split of SSv2 videos, with a focus on compositional action recognition regardless of object classes. It ensures object classes of the same action are disjoint between training and validation. The resulting split has  $\sim 55\text{K}$  training videos and  $\sim 67\text{K}$  validation videos. The videos are decoded in 12 FPS.

The Something-Else authors additionally annotate ground-truth object bounding boxes with tracked identities for videos in its split. In our analytic experiments, we use these oracle box annotations to study the performance upper-bound of our framework. In our main experiments, we use inferred bounding boxes from ORViT [21], which are from a finetuned FasterRCNN detector [37, 54] with ResNet-101 backbone on the corresponding splits of the SomethingElse box annotations.

**Epic-Kitchens** [8] contains 67K training videos and 10K validation video segments on kitchen scenes, where each segment is on average 3 seconds long. The action is composed of a verb (from 97 classes) and a noun (from 300 classes), and is considered as correct if both verb and noun classification are correct. The dataset does not provide bounding box annotations, and we follow ORViT [21] to use an off-the-shelf object detector trained on hand-object interaction dataset [40], and use a learning-free tracker SORT [4] to link them. We decode the videos at 60 FPS [21].

### 5.1. Implementation details

We implement our method in JAX [6] based on the Scenic library [10]. We use space-time ViViT [1] with VideoMAE pretraining [43] as our baseline model. In all experiments, we use input size of (16, 224, 224) and token cube size of (2, 16, 16), which results in an input size of 1568 tokens. Our training hyper-parameters follow VideoMAE finetuning [43]. Specifically, we train our model and

baselines using the AdamW optimizer [25] with learning rate  $10^{-3}$  and batch size of 256 videos with a cosine learning rate decay. During training, we resize images to short side 256 and do a random crop at 224. Unless specified, during testing we resize the image to short side 224, and do a single center crop. We use uniform frame sampling following VideoMAE [43] and TSN [46]. I.e., the sampling stride is proportional on the video length so that the sampled frames always cover the whole video. We use regularizations including drop path [22] of 0.1, label smoothing [41] of 0.1, and mixup [58] of 0.8. On SSv2, we load VideoMAE checkpoint pretrained on SSv2 training set [43], and train our model for 30 epochs. On SomethingElse and Epic Kitchens, we initialize from the VideoMAE checkpoint pretrained on Kientics400 [24] publicly released by the authors [43], and train our model for 50 epochs. All other training parameters are the same between datasets.

### 5.2. Main results

We first show our full model results with both our proposed modules, object-guided token sampling (OGS) and object-aware attention module (OAM). The hyper-parameters of our model include the foreground token split ratio  $X\%$  and the background sampling ratio  $Y\%$  (Section 4.1). We keep the background sampling ratio  $Y\% = 10\%$  (as will be ablated in Figure 7, except for 10% total tokens where we use  $Y\% = 5\%$ ), and sweep  $X\%$  so that the total number of remaining tokens  $(X + Y)\%$  range from 10% to 90%. For the results of 100% tokens, we do not apply OGS and only use OAM.

Figure 5 shows the results on SomethingElse, SSv2, and Epic-Kitchens under different number of tokens. First, we observe that without dropping tokens (x-axis at 1.0), OAM effectively boosts the space-time ViViT baseline [1] by a healthy margin, with 2.1%, 1.3%, and 0.6% points improve-




Figure 6. **Effectiveness of Object-guided sampling.** We show action classification accuracy on SomethingElse (a), SSv2 (b), and Epic Kitchens (c) dataset with our object-guided token sampling only under different numbers of tokens (without Object-aware attention). We use the same space-time ViViT baseline as Figure 5. We compare to a naive-drop baseline that uniformly random sample tokens. Our object-guided sampling consistently outperforms the naive sample baseline on all sampling ratios and datasets. Surprisingly, we observe that token-sampling alone can improve action-recognition performance on SomethingElse and SSv2 (with both inferred or oracle boxes), and matches the baseline performance with 50% ~ 60% number of tokens.

ments on the three datasets, respectively. For analysis purposes, we also show results with the ground truth detections (in both training and testing) when available. Here the improvements further increase to 7.3% and 6.3% on SomethingElse and SSv2, respectively. This implies our performance can be further improved with stronger detectors.

When starting dropping tokens using OGS, the performance of our model decreases due to dropped input information. The performance drops mildly with respect to tokens, and overall yields a favourable trade-off. Specifically, on SomethingElse, SSv2, and Epic-Kitchens, our model meets the baseline’s performance at 30%, 40%, and 60% tokens, respectively. This convincingly shows object regions are right highlights of the videos and can effectively guide token selection. Again for analysis purposes, we show when ground truth bounding boxes are available, we can match the baseline performance with only 10% tokens.

As our model saves tokens in inference, we configure a testing strategy that applies our 50%-token model with 2-view testing (i.e., uniform sampling frames with different starting frames and averaging output logits) so that the overall number of tokens processed matches the baseline. Figure 5 shows this model (★) further improves our full-token model. This again highlights that tokens-efficiency is an important aspect in action recognition.

**Effectiveness of Object-guided sampling.** Next, we apply only OGS without OAM to study the importance of the sampler alone. Following the setting in Figure 5, we vary the foreground token split ratio  $X\%$  and show the action classification change of the three datasets. We compare to a naive token dropping baseline that uniform-randomly drop tokens in space and time, which is equivalent to setting  $X\% = 0$  and varying  $Y\%$  in our framework, and is the same as the token sampler used in video MAE pretraining [15]. Figure 6 shows the results. On all the three datasets and all sampling ratios, our object-guided sampler outperforms the uniform

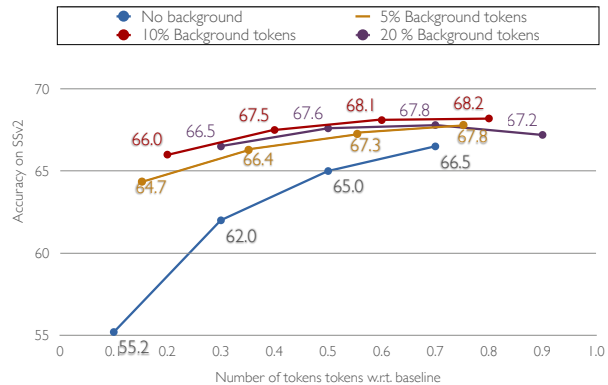


Figure 7. **Importance of adding background tokens in object-guided token sampling.** We show action recognition accuracy under different numbers of remaining tokens ( $(X + Y)\%$  in x-axis) with different amounts of background tokens (different  $Y\%$  in different lines). We experiment using ground truth boxes on Something-something V2 dataset with OGS. We observe the model without any background tokens performs poorly, and 10% background tokens give the best overall performance.

sampling baseline, with a more significant gain when the total sampling ratio is low. Surprisingly, we observe solely dropping tokens can improve action classification accuracy for free: on all the three datasets, we observe a 0.2 points gain on their corresponding optimal sampling ratio. This is likely thanks to dropping background tokens highlights the foreground motions and makes the learning task easier. This phenomenon is more pronounced when we use the oracle detections, where we observed a 2.1% gain on both Something datasets. The performance improvements do not increase with more tokens, as 100% tokens go back to the no-sampling baseline. Most importantly, on all the three datasets, we all match the full-token baseline with fewer tokens, with 60%, 50%, 90% tokens, respectively. Figure 8 shows qualitative results of our sampler.

Layer index	Top 1 Acc.	gFLOPs		Top 1 Acc.		Top 1 Acc.
*0	66.1	80	RoIAlign	67.2	w.o. tracking	67.0
1	66.0	88	Binary block mask	67.2	w. identity attention	67.4
6	65.9	130	*Center heatmap weighting	67.4	*w. identity embedding	67.4

(a) **When to drop tokens.** We apply our sampler on different layers (without object attention), with 10% background tokens and 40% foreground tokens.

(b) **Object feature aggregation functions.** We compare RoIAlign used in ORViT [21], a simpler block masking, and our proposed heatmap weighting.

(c) **How to use tracking information.** We compare no tracking, using additional layers for tokens for the same object, or simply adding an identity embedding.

Table 1. **Ablation studies.** We ablate design choices for both object-guided token sampling (1a), and object-aware attention module (1b, 1c). All models use inferred boxes on SSv2 [17]. \* means our default option.

### 5.3. Ablation studies

Next, we ablate design choices in both our proposed modules, including the importance of background tokens, when to apply token sampling, and designs in the object-aware attention module.

**Importance of adding background tokens.** We first show it is crucial to include both objects and contexts. Figure 7 shows the accuracy trade-off under different amounts of background tokens on SomethingSomegning-V2 (with ground truth bounding boxes). The entry without any background tokens significantly underperforms other entries even with very few background tokens (5%). This is because non-object tokens provide important contextual information that might be critical for the object, e.g., showing the action is happening on a desk or a dining table. Background tokens also prevent the model from overfitting into inaccurate object detection or annotation. Our experiments show including 10% background tokens gives an overall good trade-off, but different foreground-background token split ratios do not crucially impact the results.

**When to drop tokens.** OGS takes the full set of token as input and produce downsampled tokens. It can be applied to any blocks in the transformer. Intuitively, applying token downsampling in later layers [38] reserves more information with the cost of more computation. Table 1a ablates doing token sampling in different layers. We observe dropping tokens at the very beginning before any transformer layers work as well as dropping in middle blocks, while being the most computationally efficient. We thus use this early dropping by default.

**Object feature aggregation functions.** OAM aggregates patch features into object features. A popular way [21, 46, 51] to do that is to crop and resize (i.e., RoIAlign [20]) the box region in the feature grids. However, this is no longer feasible when the input features are unstructured patches after downsampling. In Table 1b, we compare alternatives to RoIAlign under the full token settings. We first use a binary block mask that masks out input features outside of the bounding box region, followed by linear layers and a global pooling. This matches RoIAlign without cropping. Our heatmap-weighted pooling (Equation 6) further improves the block-mask, likely due to the center heatmap giving a better estimation of the object mask. A mask-guided fea-

ture aggregation [7] might perform better, when instance segmentation annotations are available.

**How to use tracking information.** By default, our object-aware attention module does not use object tracking information as each object token is created per frame individually. We study options to add identity information in our framework when available (e.g., via a light-weighted tracker [4]). Table 1c shows the results. Without any tracking information, our object-aware attention module still improves the baseline (66.1) by 0.9 points. We experimented with two ways to use tracking. The first is to apply an additional attention module on object tokens from the same track, and the second is to simply add the same embedding to objects within a track following [21]. We observe both further improve the performance by 0.4 points, and we choose to use the embedding-based way for simplicity.

### 5.4. Comparisons to existing work

Next, we compare to existing methods on both Object-based video modeling and token-efficient transformers.

**Object-based video representation.** We compare our full model to ORViT [21] and ObjectLearner [57] on SomethingElse [32] and SSv2 datasets [17]. Following their evaluation protocol, we use ground truth boxes on SomethingElse and the same inferred boxes on SSv2. Both methods use MotionFormer [34] as their baseline which requires

	SSv2	S.Else	n. FLOPs
MFormer [34]	66.5	62.8	1× (370)
MFormer w. ORViT [21]	67.9 / -	69.7	1.09×
MFormer w. ObjectLearner [57]	- / 74.0	73.6	1.04×
VideoMAE (our baseline)	64.8	67.0	1× (181)
VideoMAE w. ORViT (reproduced)	68.3 / 72.8	71.6	1.06×
ObjectViViT	68.5 / 73.0	72.0	1.06×
ObjectViViT (50% tokens ×2 views)	70.0 / 74.2	74.5	1.06×

Table 2. **Compare to existing object-based video models [21, 57].** We report normalized FLOPs (n. FLOPs) with respect to the corresponding baseline. We show results on SomethingElse (S.Else) with ground truth boxes and and SSv2 using both inferred and ground truth boxes (results shown in format: inferred boxes / ground truth boxes). The results are reported with a single temporal view and 3 spatial crops [21] except for the last row, which uses 3 spatial crops and 2 temporal views with 50% tokens.

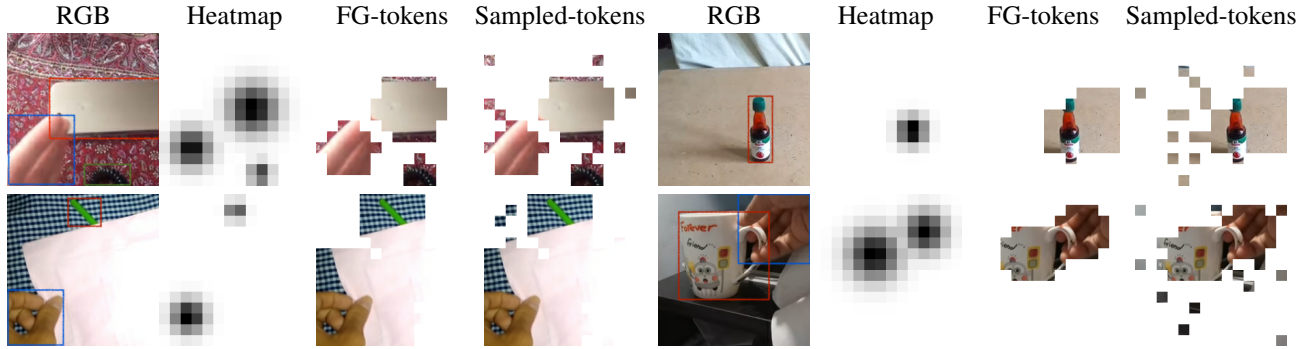


Figure 8. **Qualitative results of our object-guided token sampling.** A row shows two frame from two different videos. From left to right, we show the original RGB pixels with external bounding boxes, the object-heatmap used to label tokens, the resulting foreground tokens, and the overall sampled tokens including background. We use foreground split ratio  $X = 30\%$  and background sample ratio  $Y = 10\%$ .

full grid-shaped inputs, which are different from our vanilla ViT baseline. Thus, we reproduce ORViT [21] on our baseline based on their released code. We report their original numbers quoted from their publications as well as our reproduced numbers on our baseline. Table 2 show our full model slightly outperform ORViT [21], mostly thanks to our better object feature aggregation functions. We highlight our additional ability to process with tokens. When using our optimal model with 50% tokens and multi-view testing, our gains are larger under the same normalized FLOPs.

**Token-efficient transformers.** Finally, we put our results in the context of existing token-efficient transformers for videos. We list perform of recent work STTS [45] and AdaFocus [50], both of which adaptively select video regions for action recognition on SomethingSomething-V2 [17]. Note we are not strictly comparable to them in terms of accuracy-efficiency trade-off for two reasons. First, our model requires additional object inputs, and over-

	Ratio	Baseline	Proposed	$\Delta$
<b>End-to-end</b>				
STTS [45]	70%	69.6 [30]	68.7	-0.9%
AdaFocus v3 [50]	44%	59.6 [29]	59.6	+0.0%
<b>Object-guided</b>				
ObjectViViT	40%	66.1 [43]	66.2	+0.1%
ObjectViViT w. gt boxes	15%	66.1 [43]	67.4	+1.3%

Table 3. **Our results in contexts of other token-efficient video transformers on SomethingSomething-V2 [17].** For each method, we report numbers of tokens with respect to its corresponding baseline, the baseline accuracy, the proposed-model accuracy, and show the absolute performance change ( $\Delta$ ) w.r.t the baseline. Top block: exiting end-to-end methods. We quote the numbers of STTS [45] and AdaFocus [50] from their original publications. Bottom block: our object-guided models. Our model achieves better *token*-efficiency compared to end-to-end methods with predicted boxes from detectors, and can be further improved with ground truth boxes. Note we do not claim better *computation*-efficiency as our models require additional detection inputs.

all costs more computes when taking into account the detector. Also, different methods are developed on different video architectures, and we use performance change w.r.t. the baseline versus the remaining pixels as an evaluation proxy. Table 3 shows the results. With bounding boxes from object detectors, our result of retaining baseline performance at 40% tokens and is in a favorable position compared to STTS and AdaFocus. Our results with ground truth boxes further highlight the advantages of using objects.

## 6. Discussion

Our work is a first step towards a compact, object-based video processing framework. First, we proposed an object-guided token sampling module that could drop background regions at an aggressive ratio. Second, we proposed an object-aware attention module that utilizes object-token relation and improves action classification accuracy with minor additional computation. Our overall framework improves both the token-efficiency and classification accuracy on popular action recognition benchmarks.

Despite promising preliminary results, there are several under-explored components in our framework. First, it is still an open problem which object detector best suits the action recognition benchmark. In our work, we found domain-specific detectors trained on human-action interaction datasets [32, 40] worked the best, and switching to a stronger detection backbone [27] or a more general open-world detector [59] does not help. It is unclear how object detection metrics (*i.e.* the mAP on most datasets) correlate to the ability to help action recognition. Second, running a detector requires substantial computation, particularly since modern detectors operate on high-resolution images. As a result, our framework yet does not speed up the actual inference pipeline, and we leave integrating light-weighted object detectors within video frameworks as an exciting future direction. We hope our work opens exciting avenues in the interaction between object detection and video recognition to future video architectures more efficient and accurate.



## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 1, 2, 3, 5
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv:1607.06450*, 2016. 3
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 1, 2
- [4] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *ICIP*, 2016. 5, 7
- [5] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. *ICLR*, 2023. 3
- [6] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. 5
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022. 7
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *TPAMI*, 2020. 2, 4, 5
- [9] Mostafa Dehghani, Anurag Arnab, Lucas Beyer, Ashish Vaswani, and Yi Tay. The efficiency misnomer. In *ICLR*, 2022. 3
- [10] Mostafa Dehghani, Alexey Gritsenko, Anurag Arnab, Matthias Minderer, and Yi Tay. Scenic: A jax library for computer vision research and beyond. In *CVPR*, 2022. 5
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 2, 3
- [12] Gamaleldin F 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. *NeurIPS*, 2022. 2
- [13] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 1, 2
- [14] Mohsen Fayyaz, Soroush Abbasi Kouhpayegani, Farnoush Rezaei Jafari, Eric Sommerlade, Hamid Reza Vaezi Joze, Hamed Pirsiavash, and Juergen Gall. Ats: Adaptive token sampling for efficient vision transformers. *ECCV*, 2022. 3
- [15] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *NeurIPS*, 2022. 3, 6
- [16] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 1
- [17] 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*, 2017. 2, 4, 7, 8
- [18] Kalanit Grill-Spector and Nancy Kanwisher. Visual recognition: As soon as you know it is there, you know what it is. *Psychological Science*, 2005. 1
- [19] Abhinav Gupta and Larry S Davis. Objects in action: An approach for combining action understanding and object perception. In *CVPR*, 2007. 1
- [20] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 2, 4, 7
- [21] Roei Herzig, Elad Ben-Avraham, Karttikeya Mangalam, Amir Bar, Gal Chechik, Anna Rohrbach, Trevor Darrell, and Amir Globerson. Object-region video transformers. In *CVPR*, 2022. 1, 2, 4, 5, 7, 8
- [22] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *ECCV*, 2016. 5
- [23] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *ICLR*, 2017. 3
- [24] 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:1705.06950*, 2017. 5
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *iclr*. 2015. *ICLR*, 2015. 5
- [26] Zhenglun Kong, Peiyan Dong, Xiaolong Ma, Xin Meng, Wei Niu, Mengshu Sun, Xuan Shen, Geng Yuan, Bin Ren, Hao Tang, et al. Spvit: Enabling faster vision transformers via latency-aware soft token pruning. In *ECCV*, 2022. 3
- [27] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *ECCV*, 2022. 8
- [28] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvitv2: Improved multiscale vision transformers for classification and detection. In *CVPR*, 2022. 2
- [29] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 8
- [30] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *CVPR*, 2022. 1, 2, 8
- [31] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *NeurIPS*, 2020. 2
- [32] Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, and Trevor Darrell. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *CVPR*, 2020. 2, 3, 4, 5, 7, 8

- [33] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *CVPR*, 2022. 3
- [34] Mandela Patrick, Dylan Campbell, Yuki Asano, Ishan Misra, Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and João F Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. *NeurIPS*, 2021. 2, 4, 7
- [35] Gorjan Radevski, Marie-Francine Moens, and Tinne Tuytelaars. Revisiting spatio-temporal layouts for compositional action recognition. *BMVC*, 2021. 1, 2
- [36] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *NeurIPS*, 2021. 1, 3
- [37] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*, 2015. 3, 5
- [38] Michael S Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. Tokenlearner: What can 8 learned tokens do for images and videos? *NeurIPS*, 2021. 7
- [39] Aviv Shamsian, Ofri Kleinfeld, Amir Globerson, and Gal Chechik. Learning object permanence from video. In *ECCV*, 2020. 2
- [40] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F Fouhey. Understanding human hands in contact at internet scale. In *CVPR*, 2020. 5, 8
- [41] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016. 5
- [42] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *science*, 2011. 1
- [43] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NeurIPS*, 2022. 3, 5, 8
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 2, 3
- [45] Junke Wang, Xitong Yang, Hengduo Li, Zuxuan Wu, and Yu-Gang Jiang. Efficient video transformers with spatial-temporal token selection. *ECCV*, 2022. 1, 3, 8
- [46] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *TPAMI*, 2018. 5, 7
- [47] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018. 1, 2
- [48] Yulin Wang, Zhaoxi Chen, Haojun Jiang, Shiji Song, Yizeng Han, and Gao Huang. Adaptive focus for efficient video recognition. In *ICCV*, 2021. 1, 3
- [49] Yulin Wang, Yang Yue, Yuanze Lin, Haojun Jiang, Zihang Lai, Victor Kulikov, Nikita Orlov, Humphrey Shi, and Gao Huang. Adafocus v2: End-to-end training of spatial dynamic networks for video recognition. In *CVPR*, 2022. 3
- [50] Yulin Wang, Yang Yue, Xinhong Xu, Ali Hassani, Victor Kulikov, Nikita Orlov, Shiji Song, Humphrey Shi, and Gao Huang. Adafocusv3: On unified spatial-temporal dynamic video recognition. In *ECCV*, 2022. 1, 3, 8
- [51] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019. 2, 7
- [52] Chao-Yuan Wu and Philipp Krahenbuhl. Towards long-form video understanding. In *CVPR*, 2021. 2
- [53] Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R Manmatha, Alexander J Smola, and Philipp Krähenbühl. Compressed video action recognition. In *CVPR*, 2018. 1
- [54] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 5
- [55] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview transformers for video recognition. In *CVPR*, 2022. 1
- [56] Hongxu Yin, Arash Vahdat, Jose M Alvarez, Arun Mallya, Jan Kautz, and Pavlo Molchanov. A-vit: Adaptive tokens for efficient vision transformer. In *CVPR*, 2022. 3
- [57] Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Is an object-centric video representation beneficial for transfer? *ACCV*, 2022. 2, 7
- [58] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ICLR*, 2018. 5
- [59] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Phillip Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. *ECCV*, 2022. 8
- [60] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv:1904.07850*, 2019. 3, 4