

# Few-Shot Video Object Detection

Qi Fan<sup>1</sup>, Chi-Keung Tang<sup>1</sup>, and Yu-Wing Tai<sup>1,2</sup>

<sup>1</sup> The Hong Kong University of Science and Technology, <sup>2</sup> Kuaishou Technology  
fanqics@gmail.com, cktang@cs.ust.hk, yuwing@gmail.com

**Abstract.** We introduce Few-Shot Video Object Detection (FSVOD) with three contributions to real-world visual learning challenge in our highly diverse and dynamic world: 1) a large-scale video dataset FSVOD-500 comprising of 500 classes with class-balanced videos in each category for few-shot learning; 2) a novel Tube Proposal Network (TPN) to generate high-quality video tube proposals for aggregating feature representation for the target video object which can be highly dynamic; 3) a strategically improved Temporal Matching Network (TMN+) for matching representative query tube features with better discriminative ability thus achieving higher diversity. Our TPN and TMN+ are jointly and end-to-end trained. Extensive experiments demonstrate that our method produces significantly better detection results on two few-shot video object detection datasets compared to image-based methods and other naive video-based extensions. Codes and datasets are released at <https://github.com/fanq15/FewX>.

**Keywords:** few-shot video object detection, object indexing/retrieval, tube proposal network, temporal matching network

## 1 Introduction

We ask the following question: *Given a bunch of videos, how can we index and localize all novel objects of interest as video clips?* See Figure 1.

This problem is becoming increasingly essential with massive video collections in this media era: movies, YouTube videos, TikTok streaming videos, surveillance videos, just to name a few. The video objects of interests can be highly novel, often personalized, and thus are not covered by any existing datasets. Marvel fans may want to collect all Iron Man or Hulk clips from all Marvel movies, while warfare collectors want to create a TikTok video consisting of tank clip collections from war movies. We may not even know which videos contain the interested objects.

No existing tasks or solutions can solve this real-world challenge. Notably, multiple object tracking [13,90], image/video object detection [36,96,3,88] are all restricted in fixed and limited training classes. Single object tracking [5,67] can track new classes, but it requires user-provided template for every video and can

---

This research was supported in part by Kuaishou Technology, and the Research Grant Council of the Hong Kong SAR under grant No. 16201420.



**Fig. 1.** Given only a few support objects of interest, our FSVOD detects all objects of the same category in query videos. Note that FSVOD enables object indexing/retrieval in a bunch of query videos to extract video clips containing the target objects.

only track the target template object. Few-shot learning seems a good candidate solution. But existing few-shot object detection [134,31] and few-shot classification [58,114] are specifically designed for still images and they will produce numerous false positive results in videos. Few-shot video classification [10,146,55] does not target at instance recognition.

This real-world challenge motivates few-shot video object detection (FSVOD): given only a few *support images* of the target object in an unseen class, FSVOD detects all the objects belonging to the same class in a given *query video*. The given support images can be arbitrary objects of interest, and FSVOD works on arbitrary videos for indexing and localization. The key to successful FSVOD is simultaneously modeling both high dynamics and high diversity of our dynamic and diverse world, while other existing tasks can only contribute either high dynamic or high diversity, as summarized in Table 1, and thus falling short of the real-world challenge.

The technical contributions of FSVOD, namely, Temporal Proposal Network (TPN) for high object dynamics and Temporal Matching Network (TMN+) for high object diversity, will be detailed. The core idea is to perform temporal matching between the tube-aggregated query features and supports, which enables high-quality detection based on the representative tube features and eliminates ghost objects (false positive predictions) which heavily suffers the few-shot image object detection methods.

The other contribution of this paper consists of a large-scale dataset that enables new research on few-shot video object detection. Our dataset contains 500 classes with a small and balanced number of high-quality videos in each class. The numerous classes with class-balanced videos enable the trained model to learn a general relation metric for novel classes. Note that this dataset contributes not only as the first benchmark for FSVOD, but also as a useful benchmark for other important vision tasks, such as multi-object tracking and video

**Table 1. Comparing FSVOD and relevant computer vision tasks** in terms of *dynamic* and *diversity* capabilities: detecting *box* for *novel* object classes and/or *multiple* objects, and whether *temporal* information is considered. Number of ‘+’ indicates how diversity each task can contribute. ‘S.A.’ means scene adaptation.

Task	Dyn.	Div.	Box	Nov.	Mul.	Temp.	S.A.
Image Object Detection (IOD)	no	+	✓	✗	✓	✗	✓
Video Object Detection (VOD)	yes	+	✓	✗	✓	✓	✓
Multiple Object Tracking (MOT)	yes	+	✓	✗	✓	✓	✗
Single Object Tracking (SOT)	yes	+	✓	✓	✗	✓	✗
Few-Shot Classification (FSC)	no	++	✗	✓	✓	✗	✓
Few-Shot Object Detection (FSOD)	no	+++	✓	✓	✓	✗	-✓
<b>Few-Shot Video Object Detection</b>	yes	+++	✓	✓	✓	✓	✓

object detection which are still in lack of a well-constructed, class-balanced video benchmark on par in the number of classes as FSVOD-500.

## 2 Related Work

The FSVOD task is related to few-shot learning, object detection and video understanding. Table 1 summarizes its relationship with closely related tasks.

**Few-Shot Classification (FSC).** Optimization-based works learn task-agnostic knowledge on model parameters [35,4,65] for fast adaptation to new tasks on limited training data, using only a few gradient update steps. Some works [42,117] hallucinate new images for novel classes from limited labeled data. Metric-based methods exploit a weight-shared network [58] to extract features of the support and query images before feeding them to a transferable distance metric. Such matching strategy [114,137,130,104] captures inherent variety between supports and queries irrespective of classes and thus can be directly applied for classifying novel classes.

**Few-Shot Object Detection (FSOD).** With encouraging progress made in the few-shot classification, few-shot learning has continued to contribute to important computer vision tasks [26,84,48,39,79,70,30] at a fast pace especially for object detection [134,92,118]. In LSTD [11] the gap between the source and target domain is minimized. RepMet [53] learns the multi-modal distribution of the training classes in the embedding space. FR [50] exploits a meta feature learner to quickly adapt to novel classes. Some works exploit semantic relation reasoning [144], restore negative information [132], feature hallucination [140] or other techniques [105,47,138,72,66,32,69,129,116,122,126] to facilitate few-shot object detection. All of the above methods however require fine-tuning on novel classes. In FSOD [31] the authors proposed to learn a matching metric with attention RPN and multi-relation detector to detect novel classes.

Our FSVOD extends FSOD task to the temporal domain, with the technical approach motivated by the matching network [114] and FSOD network [31] to detect novel classes without fine-tuning.

**Image Object Detection (IOD).** Existing object detection methods can be mainly categorized to the two-stage approach [36,96,73] and one-stage ap-

proach [74,94,95,77,76,139], based on whether a region-of-interest proposal step is used. The two-stage approach was pioneered by R-CNN [37]. In recent years, this approach has been improved by various excellent works and achieved remarkable performance [43,103,8,9,14,71,1]. The one-stage approach on the other hand discards the proposal generation procedure in lieu of higher computational efficiency and faster inference speed with anchor-based [60,101,147,141] or anchor-free detectors [64,80,143,27,133,111,78,59].

**Video Object Detection (VOD).** Video object detection aims at detecting objects of pre-defined classes in a given video. Some enhance the quality of per-frame features by integrating temporal information locally [3,24,124,115], globally [22,102,121] or both [88,119,120,127,12], while others follow the “sequential detection tracking” paradigm [34,148,149,150,51,107] to associate and rescore detected boxes on individual frames. The above work in intensive supervision and cannot be applied readily to detect novel classes. VOD variants include e.g., video object segmentation (VOS) [91,128], video instance segmentation (VIS) [131] and video panoptic segmentation (VPS) [57].

Both IOD and VOD are restricted to pre-defined classes making it hard for them to detect novel classes. FSVOD eliminates this restriction with its detection generality on novel classes in videos.

**Single Object Tracking (SOT).** Given an arbitrary target with its location in the first frame, single object tracking aims to infer its location in subsequent frames of the given video. Thanks to the construction of new benchmark datasets [29,123] and annually held tracking challenges [63,61,62], we have witnessed rapid performance boost in the last decade. The correlation filter based trackers [17,19,20,46] achieve superb performance with efficient inference speed. The recent emerging siamese network based trackers [5,40,68,67,45,109,112] have drawn much attention due to their well-balanced performance and efficiency.

Although SOT models can track unseen objects, they heavily rely on the provided template and can only track one target object. The online tracking trackers [6,38,18,7,20,17,19,16,83] can be finetuned/updated on the first frame, but they focus on tracking single object with the video-specific annotated first frame. On other hand, our FSVOD focuses on detecting arbitrary novel objects in videos based on given video-agnostic support images even from other images/videos and can be reused for all input videos.

**Multiple Object Tracking (MOT).** This task [113,82] requires simultaneous prediction of spatio-temporal location and classification of video objects into pre-defined classes. Current mainstream trackers [2,81,13,136,142,56,97,108,33,135] adopt tracking-by-detection (TBD) by first performing per-frame detection and then associating the detected boxes in the temporal dimension. Some works leverage trajectories or tubes to capture motion trails of targets [90,52,100,145,89].

While MOT models can simultaneously track multiple objects, they cannot generalize to novel classes. FSVOD can detect novel classes in videos. Our technical approach is inspired by these previous methods, especially tube-based MOT and VOD methods, e.g., CPN [107] and CTracker [90], which are restricted in limited training classes.

### 3 Proposed Method

Few-shot video object detection aims at detecting *novel* classes unseen in the training set. Given a *support image* containing one object of the support class  $c$  and a *query video* sequence with  $T$  frames, the task is to detect all the objects belonging to the support class  $c$  in every frame. Suppose the support set contains  $N$  classes with  $K$  samples for each class, the problem is defined as  $N$ -way  $K$ -shot detection. Specifically, during inference, if all the support classes are exploited for detection, it is dubbed full-way evaluation.

#### 3.1 Overview

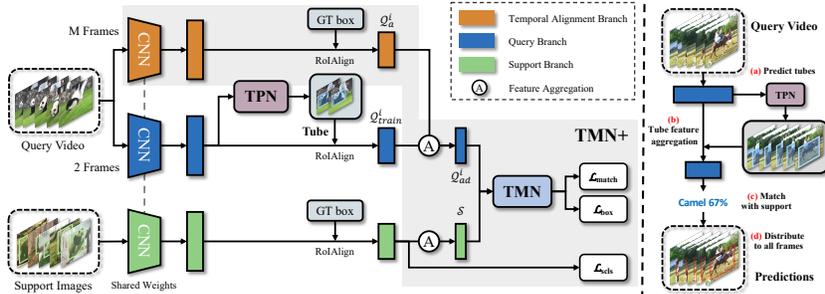
Technically, it is non-trivial to transfer few-shot learning [58,98,35] to the video object detection domain for simultaneously modeling the dynamic and diverse world. Few-shot learning requires a large-scale, class-balanced dataset with numerous base classes to train a class-agnostic metric capable of generalizing to novel classes [99,31,70]. Besides, videos present additional data challenges caused by *e.g.*, motion blur, occlusion and deformation of objects, making infeasible straightforward extension of few-shot image to few-shot video object detection without adequate temporal consideration.

This paper extends the traditional video object detection to detect novel classes in a few-shot learning setting which is not a straightforward problem. We propose a novel tube-based few-shot video object detection model for detecting novel classes in a given video, without any fine-tuning or retraining. We make the following contributions:

We first model *dynamic objects* by generating temporal tubes using our novel Tube Proposal Network (TPN) exploiting spatial adjacency and appearance similarity in the neighboring frames. Specifically, by introducing novel inter-frame proposals to detect objects in consecutive frames, TPN can capture potential objects in the query video while filtering out background and ghost objects (the false positive objects detected in isolated frames). We argue that the aggregated features across frames can better represent the target objects which leads to significant improvement on the detection performance.

Then we model *diversity of objects* using subsequent Temporal Matching Network (TMN+), which is specially designed and strategically improved to match support features and the aggregated query features from temporal tube proposals generated by TPN. Our proposed TMN+ effectively leverages the representative tube features by bridging the gap between training and inference via our novel temporal alignment branch. Furthermore, a new support classification loss is used to learn a highly discriminative feature, and a label-smoothing regularization is used for better generalization on novel unseen classes. Consequently, our TMN+ boosts matching performance on novel classes without extra computation overhead at inference.

The TPN and TMN+ are integrated into one unified network and jointly optimized in an end-to-end manner to simultaneously handle high dynamics and diversity in visual object detection.



**Fig. 2. Network architecture** at training (left) and testing (right) stages. The query video and support images are processed by the weight-shared backbone. The query video branch only processes two query images. The temporal alignment branch (TAB) is used for query feature alignment, and a classification module is introduced to produce representative support features. For clarity we show the detection on a single object, while our model can perform multi-object detection with corresponding tubes.

### 3.2 Few-Shot Video Object Detection Network

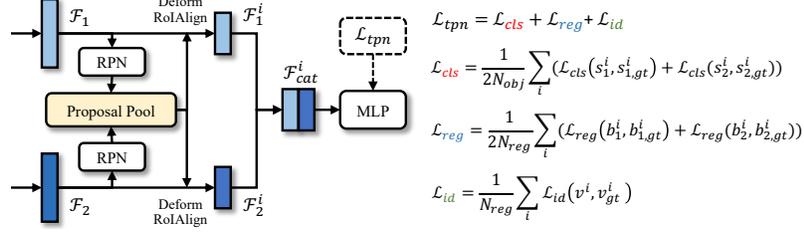
Figure 2 shows the network architecture. We propose a novel temporal detection network that exploits tubes to locate and represent objects in the temporal domain, which are then matched with support features.

**Tube Proposal Network** In image object detection, region proposal network RPN [96] has become a classical module to produce proposals for potential objects while filtering out the background. These proposals are fed to the R-CNN head for finer classification and localization.

We extend RPN to the temporal domain to generate *tube proposals* to locate and represent objects across frames. The resulting network is our novel tube proposal network (Figure 3) which exploits the high likelihood that the same object in neighboring frames tend to have *similar location and appearance*.

To utilize the location cue in adjacent frames, we propose the novel *inter-frame proposals* by feeding the same proposals to two adjacent frames. Note that proposals usually serve as a coarse prediction prior for later finer regression. The predicted boxes regressed from the same proposals indicate the same objects and therefore inter-frame proposals can associate objects across frames. However, it is also possible that objects with large motion may locate far away in adjacent frames, or the locations are occupied by other objects in the next frame. To address this problem, we adopt the deformable RoIAlign [15] operator to enlarge the search region for the target objects by adapting the sampling bins conditioned on the input feature. To exploit the appearance cue in neighboring frames to address the second problem, we verify the same object by predicting the identification score of the predicted boxes regressed from the same proposal.

Specifically, given two adjacent frames  $\{I_1, I_2\}$ , we first use RPN to generate proposals for each frame and collect both frame proposals to construct the proposal pool. Each proposal  $p_i$  in the proposal pool is simultaneously fed to the two frames to extract proposal features  $\{\mathcal{F}_1^i, \mathcal{F}_2^i\}$  with the deformable RoIAlign operator. These proposal features from individual frames are concatenated as



**Fig. 3.** Tube Proposal Network (TPN) and the loss function. The  $*_{gt}$  is the ground-truth label for the corresponding prediction,  $\mathcal{L}_{cls}$  and  $\mathcal{L}_{id}$  are both cross-entropy loss and  $\mathcal{L}_{reg}$  is the smooth  $\mathcal{L}_1$  loss.  $N_{obj}$  and  $N_{reg}$  are respectively the number of proposals and foreground proposals.

$\mathcal{F}_{cat}^i = \text{concat}(\mathcal{F}_1^i, \mathcal{F}_2^i)$ , which is then fed to the following multilayer perceptron (MLP) layer to perform objectness classification  $\{s_1^i, s_2^i\}$ , box regression  $\{b_1^i, b_2^i\}$  for each frame, and identify verification  $v^i$ . The 2-frame tube prediction is trained with the TPN loss  $\mathcal{L}_{tpn}$ , as shown in Figure 3.

During inference, the TPN needs to connect all frames in the given video by repeating the 2-frame tube prediction. Consider the 3-frame case where the  $T$ -frame ( $T > 3$ ) can be generalized<sup>1</sup>: given  $\{I_1, I_2, I_3\}$ , we first send  $\{I_1, I_2\}$  to the model to generate a 2-frame tube  $\{b_1, b_2\}$ . Then we feed the pre-computed tube box  $b_2$  to  $\{I_2, I_3\}$  as the inter-frame proposal to generate the tube box  $b_3$  for frame  $I_3$  to construct another 2-frame tube  $\{b_2, b_3\}$ . We can construct a 3-frame tube  $\{b_1, b_2, b_3\}$  by linking  $\{b_1, b_2\}$  and  $\{b_2, b_3\}$  through the inter-frame proposal  $b_2$ . The overlapped frame  $I_2$  is used to verify the same objects between two frame pairs and its feature are reused in the process to avoid repeating computation as in CTracker [90]. Thus, we can sequentially detect tube boxes for all the frames and generate tube proposals.

**Temporal Matching Network** After obtaining tube proposals, we extract and aggregate tube features and compare them with support features using a matching network, where the matching results are then distributed to the tube proposals in all frames. We re-design the matching network (MN) in the temporal domain to take advantage of tube features. Consequently, our discriminative temporal matching network TMN+ and TPN which share backbone features are jointly trained for better optimization. Below we detail the design rationale on a single instance/track  $i$ , starting from MN, TMN and finally TMN+, and it is easy to apply them on multiple objects of different classes.

**MN.** From  $\{I_1, I_2\}$ , the **query branch** of backbone extracts query features  $\{\mathcal{Q}_1^i, \mathcal{Q}_2^i\}$  for each proposal  $p_i$  of instance  $i$  with RoIAlign operator. The **support branch** extracts the support features  $\mathcal{S}$  in the ground-truth boxes of the support images. The MN then computes the distance between  $\mathcal{Q} = \frac{1}{2}(\mathcal{Q}_1^i + \mathcal{Q}_2^i)$  and  $\mathcal{S}$  and classifies  $\mathcal{Q}$  to the nearest support neighbor. We adopt the multi-relation head with contrastive training strategy from FSOD [31] as our matching network

<sup>1</sup> The operations are parallel conducted for each instance/track  $i$  and we omit the instance notion for simplicity.

(MN) for its high discriminative power. Refer to the supplementary material for more details about its architecture.

**TMN.** The above MN is however designed for image object detection and is unsuitable to be applied in the temporal domain. The main problem is the misalignment between training and inference for the query features  $\mathcal{Q}$ : In the training stage,  $\mathcal{Q}_{train} = \frac{1}{2}(\mathcal{Q}_1^i + \mathcal{Q}_2^i)$  only involves the proposal feature in two frames, limited by the GPU memory and the joint training with TPN. While in the inference stage,  $\mathcal{Q}_{test} = \frac{1}{T}(\mathcal{Q}_1^i + \mathcal{Q}_2^i + \dots + \mathcal{Q}_T^i)$  is derived from all the frames in the tube proposal. This misalignment can produce bad matching result and overall performance degradation.

To bridge this training and inference gap, we propose a novel temporal matching network (TMN) by introducing a *temporal alignment branch (TAB)* for query feature alignment. Specifically, for proposal  $p_i$  of the target object  $i$ ,  $\mathcal{Q}_{train} = \frac{1}{2}(\mathcal{Q}_1^i + \mathcal{Q}_2^i)$  involves two frames  $\{I_1, I_2\}$ , and the TAB randomly selects images<sup>2</sup> from remaining frames  $\{I_3, I_4, \dots, I_T\}$  and extracts the aligning features  $\mathcal{Q}_a = \frac{1}{M}(\mathcal{Q}_3^i + \mathcal{Q}_4^i + \dots + \mathcal{Q}_M^i)$  for the target object  $i$ , where  $M$  is the number of selected aligning query images. Then we generate the aligned query feature  $\mathcal{Q}_{ad}^i = \alpha\mathcal{Q}_{train}^i + (1 - \alpha)\mathcal{Q}_a^i$  as the feature aggregation to represent the target object and perform matching with supports in the training stage. Our TMN thus bridges this gap without disrupting the design of TPN and without introducing additional computational overhead by removing TAB at inference time.

The loss function is  $\mathcal{L}_{tmn} = \mathcal{L}_{match} + \mathcal{L}_{box}$ , where  $\mathcal{L}_{match}$  is the cross-entropy loss for binary matching and  $\mathcal{L}_{box}$  is the smooth  $\mathcal{L}_1$  loss for box regression.

**TMN+.** To enhance discriminative ability, TMN+ incorporates label-smoothing regularization [106] into TMN for better generalization and a jointly optimized support classification module for more representative feature.

We first introduce label smoothing to the matching loss  $\mathcal{L}_{match}$  of TMN, which is widely used to prevent overfitting in the classification task [87,110] by changing the ground-truth label  $y_i$  to  $y_i^* = (1 - \varepsilon)y_i + \frac{\varepsilon}{\beta}$ , where  $\varepsilon$  is the constant smoothing parameter and  $\beta$  is the number of classes. This prevents the model from being overconfident to the training classes and is therefore inherently suitable for the few-shot learning models focusing on the generalization on novel classes. Then, we add a support classification module (classifier) to the support branch to enhance the intra-class compactness and inter-class separability in the Euclidean space and thus generate more representative features for matching in TMN. We adopt cross-entropy loss as its loss function  $\mathcal{L}_{scls}$ .

During training, the TPN and TMN+ are jointly and end-to-end optimized with the weight-shared backbone network by integrating all the aforementioned loss functions:

$$\mathcal{L} = \lambda_1\mathcal{L}_{tpn} + \lambda_2\mathcal{L}_{tmn} + \lambda_3\mathcal{L}_{scls} \quad (1)$$

where  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are hyper-parameter weights to balance the loss functions and are set to 1 in our experiments.

<sup>2</sup> The random selection can be regarded as data augmentation to imitate the imperfect tube features during inference.

## 4 FSVOD-500 Dataset

There exist a number of public datasets with box-level annotations for different video tasks: ImageNet-VID [23] for video object detection; LaSOT [29], GOT-10k [49], Youtube-BB [93], and TrackingNet [86] for single object tracking; MOT [85], TAO [21], Youtube-VOS [128] and Youtube-VIS [131] for multi-object tracking. However, none of these datasets meet the requirement of our proposed few-shot video object detection task. Some datasets (Youtube-BB [93], TrackingNet [86], ImageNet-VID [23], Youtube-VOS [128] and Youtube-VIS [131]) have many videos but limited classes, whereas a sufficiently large number of base classes is essential to few-shot learning. On the other hand, although other datasets (GOT-10k [49] and TAO [29]) contain diverse classes, not all instances of the same target class are annotated in a video, and therefore are not suitable for the few-shot task. Last but not least, all of these datasets are not specifically designed for few-shot learning whose train/test/val sets are class-overlapping and cannot be used to evaluate the generality on unseen classes.

Thus, we design and construct a new dataset for the development and evaluation of few-shot video object detection task. The design criteria are:

- The dataset should consist of *highly-diversified classes* for learning a general relation metric for novel classes.
- The dataset should be *class-balanced* where each class has similar number of samples to avoid overfitting to any classes, given the long-tailed distribution of many novel classes in the real world [41].
- The train/test/val sets should contain *disjoint* classes to evaluate the generality of models on novel classes.

To save human annotation effort as much as possible, rather than building our dataset from scratch, we exploit existing large-scale video datasets for supervised learning, *i.e.*, LaSOT [29], GOT-10k [49], and TAO [21] to construct our dataset subject to the above three criteria. The dataset construction pipeline is consist of dataset filtering, balancing and splitting.

**Dataset Filtering.** Note that the above datasets cannot be directly used since they are only partially annotated for tracking task: although multiple objects of a given class are present in the video, only some or as few as one of them is annotated while the others are not annotated. Thus, we filter out videos with non-exhaustive labels while keeping those with high-quality labels covering all objects in the same class (target class). We also remove videos containing extremely small objects which are usually in bad visual quality and thus unsuitable for few-shot learning. Note that exhaustive annotation for all possible classes in such a large dataset is expensive and infeasible [21,41]. Therefore, only the target classes are exhaustively annotated for each video while non-target classes are categorically ignored.

**Dataset Balancing.** It is essential to maintain good data balancing in the few-shot learning dataset, so that sufficient generality to novel classes can be achieved without overfitting to any dominating training classes. Thus, we remove ‘person’ and ‘human face’ from the dataset which are in massive quantities (and they have

**Table 2. Dataset statistics** of FSVOD-500 and FSYTV-40. ‘‘Class Overlap’’ denotes the class overlap with MS COCO [75] dataset.

	FSVOD-500			FSYTV-40	
	Train	Val	Test	Train	Test
label FPS	1	1	1	6	6
# Class	320	80	100	30	10
# Video	2553	770	949	1627	608
# Track	2848	793	1022	2777	902
# Frame	60432	14422	21755	41986	19843
# Box	65462	15031	24002	66601	27924
Class Overlap	Yes	No	No	Yes	No
Exhaustive	Only target classes			All classes	

already been extensively studied in many works and tasks [28,25,125,54]). Then, we manually remove easy samples for those classes with more than 30 samples. Finally, each class in our dataset has at least 3 videos and no more than 30.

**Dataset Splitting.** We summarize a four-level label system (shown in the supplementary material) to merge these datasets by grouping their leaf labels with the same semantics (*e.g.*, truck and lorry) into one class. Then, we select third-level node classes similar to the COCO [75] classes and exploit their leaf node as the training classes. The remaining classes are very distinct from COCO classes, and thus used to construct the test/val sets by randomly splitting the node classes. In this way, we can take advantage of the pre-training model on COCO dataset, while the test/val classes are rare novel classes and thus complying to the few-shot setting. We follow three guidelines for dataset split:

**G1:** The split should be in line with the few-shot learning setting, *i.e.*, train set should contain common classes in the real world, while test/val sets should contain rare classes. **G2:** To take advantage of pre-training on other datasets, the train set should have a large overlap with existing datasets while the test/val sets should have largely no overlap. **G3:** The train and test/val sets should have different node classes to evaluate the generality on novel classes in a challenging setting to avoid the influence of similar classes across sets, *e.g.*, if the train set has ‘Golden Retriever’, it is much easier to detect ‘Labrador Retriever’ in the test set, which is undesirable.

Consequently, **FSVOD-500** is the first benchmark specially designed for few-shot video object detection in evaluating the performance of a given model on novel classes.

## 5 Experiments

We conduct extensive experiments to validate the effectiveness of our proposed approach. Since this is the first paper on FSVOD, we compare with state-of-the-art (SOTA) methods of related tasks by adapting them to the FSVOD task.

**Training.** Our model is trained on four GeForce GTX 1080Ti GPUs using the SGD optimizer with 45,000 iterations. The initial learning rate is set to 0.002 which decays by a factor of 10 respectively in 30,000 and 40,000 itera-

tions. Each GPU contains five cropped support images, two query images and  $M$  cropped aligning query images in the same video, where  $M$  is randomly sampled from  $[1, 10]$ . We use ResNet50 [44] as our backbone which is pre-trained on ImageNet [23] and MS COCO [75]<sup>3</sup> for stable low-level features extraction and better convergence. The model is trained with 2-way 5-shot contrastive training strategy proposed in FSOD [31]. Other hyper-parameters are set as  $\alpha = 0.5, \varepsilon = 0.2, \beta = 2$  in our experiments.

**Evaluation.** We adopt the full-way 5-shot evaluation (exploit all classes in the test/val set with 5 images per class as supports for evaluation) in our experiments with standard object detection evaluation metrics, *i.e.*,  $AP$ ,  $AP_{50}$ , and  $AP_{75}$ . The evaluations are conducted 5 times on randomly sampled support sets and the mean and standard deviation are reported. Refer to the supplemental material for more training and evaluation details.

**FSYTV-40.** To validate model generalization on datasets with different characteristics, we construct another dataset built on Youtube-VIS dataset [131] for the FSVOD task. FSYTV-40 is vastly different from FSVOD-500 with only 40 classes (30/10 train/test class split following the same dataset split guidelines above, with instances of all classes are exhaustively annotated in each video), more videos in each class and more objects in each video. Table 4 tabulates the detailed statistics of both datasets.

## 5.1 Comparison with Other Methods

With no recognized previous work on FSVOD, we adapt representative models from related tasks to perform FSVOD, such as image object detection (Faster R-CNN [96], and FSOD [31]), video object detection (MEGA [12] and RDN [24]) and multiple object tracking (CTracker [90], and FairMOT [136], and CenterTrack [142]). Only FSOD model can be directly applied frame-by-frame to perform FSVOD. For others, we exploit their models to generate class-agnostic boxes and adopt the multi-relation head trained in the FSOD [31] model to evaluate the distance between the query boxes and supports. We first perform comparison on FSVOD-500, and then generalize to FSYTV-40 (Table 3).

**Comparison with IOD-based methods.** FSOD serves as a strong baseline with its high recall of attention-RPN and powerful generalization of multi-relation head. With the same matching network, Faster R-CNN produces inferior performance due to the lower recall of its generated boxes. With the representative aggregated query feature from TPN and discriminative TMN+ in the temporal domain, our FSVOD model outperforms FSOD by a large margin.

**Comparison with VOD-based methods.** VOD-based methods operate similarly to IOD-based methods in its per-frame object detection followed by matching with supports and thus both suffer from noisy proposals and less powerful features. Interestingly, we find that VOD-based methods have a worse performance because they produce excessive proposals which heavily burden the subsequent matching procedure despite their higher recalls.

<sup>3</sup> There is no overlap between MS COCO and the val/test sets of both FSVOD-500 and FSYTV-40 datasets.

**Table 3. Experimental results on FSVOD-500 and FSYTV-40 test set for novel classes with the full-way 5-shot evaluation.**

Method	Tube	FSVOD-500			FSYTV-40		
		AP	AP <sub>50</sub>	AP <sub>75</sub>	AP	AP <sub>50</sub>	AP <sub>75</sub>
FR-CNN [96]	✗	18.2±0.4	26.4±0.4	19.6±0.5	9.3±1.4	15.4±1.7	9.6±1.7
FSOD [31]	✗	21.1±0.6	31.3±0.5	22.6±0.7	12.5±1.4	20.9±1.8	13.0±1.5
MEGA [12]	✗	16.8±0.3	26.4±0.5	17.7±0.3	7.8±1.1	13.0±1.9	8.3±1.1
RDN [24]	✗	18.2±0.4	27.9±0.4	19.7±0.5	8.1±1.1	13.4±2.0	8.6±1.1
CTracker [90]	✓	20.1±0.4	30.6±0.7	21.0±0.8	8.9±1.4	14.4±2.5	9.1±1.3
FairMOT [136]	✓	20.3±0.6	31.0±1.0	21.2±0.8	9.6±1.6	16.0±2.2	9.5±1.4
CenterTrack [142]	✓	20.6±0.4	30.5±0.9	21.9±0.4	9.5±1.6	15.6±2.0	9.7±1.3
Ours	✓	<b>25.1±0.4</b>	<b>36.8±0.5</b>	<b>26.2±0.7</b>	<b>14.6±1.6</b>	<b>21.9±2.0</b>	<b>16.1±2.1</b>

**Comparison with MOT-based methods.** MOT-based methods have a similar detection mechanism to our approach, by first generating tubes for query objects and representing them with the aggregated tube features, followed by matching between query tube features and support features. Thus, even with much lower recalls ( $\sim 70.0\%$  *v.s.*  $\sim 80.0\%$ ), they still have better performance than VOD-based methods by taking advantage of temporal matching. However, our approach still outperforms MOT-based methods by a significant margin leveraging our jointly optimized TPN and TMN+ with more representative features and powerful matching network.

**Generalization on FSYTV-40 dataset.** This dataset is very different from FSVOD-500 with the former having significantly less classes but more videos in each class, more tracks in each video, and higher annotation FPS. Although our method still outperforms other methods on this dataset, a substantial performance degradation in comparison with FSVOD-500 is resulted, which is caused by the much reduced class diversity for the matching network to learn a general relation metric for novel classes. To verify this, we train our model on the FSVOD-500 train set and evaluate it on the FSYTV-40 test set<sup>4</sup>. It can promote the performance from 14.6 to 17.8 AP. The resulting large performance boost again validates the importance of high diversity of training classes, one of the desirable properties of our FSVOD-500 for few-shot video object learning.

## 5.2 Ablation Studies

Table 4 tabulates the ablation studies on the proposal box generation network and matching (classification) network. Compared to RPN, our proposed TPN improves the performance by 3.7 AP with the same matching network. Although RPN and TPN have similar recall performance (76.2% vs 76.8%), TPN has a better classification performance due to its discriminative and aggregated temporal features, and therefore producing better detection and matching performance.

For the matching network, the RN (Relation Network [130]) based baseline performs worst which is limited by its weak matching ability. Replacing RN

<sup>4</sup> There is no overlapping or similar classes between them.

**Table 4. Ablation experimental results on FSVOD-500 val set** for 80 novel classes with the full-way 5-shot evaluation. “LSR” denotes label-smoothing regularization and “SCM” denotes support classification module.

Box	Matching	$AP$	$AP_{50}$	$AP_{75}$
RPN	RN	10.1 $\pm$ 0.5	14.0 $\pm$ 0.6	11.1 $\pm$ 0.7
	MN	19.5 $\pm$ 0.9	27.4 $\pm$ 1.2	21.8 $\pm$ 1.1
TPN	MN	23.2 $\pm$ 1.2	32.7 $\pm$ 1.5	25.6 $\pm$ 1.5
	TMN	26.4 $\pm$ 1.5	37.2 $\pm$ 1.4	29.5 $\pm$ 1.6
	TMN w/ LSR	27.9 $\pm$ 1.3	39.6 $\pm$ 1.2	30.8 $\pm$ 1.5
	TMN w/ SCM	29.4 $\pm$ 0.8	41.8 $\pm$ 1.1	31.9 $\pm$ 1.2
	TMN+	<b>30.0</b> $\pm$ 0.8	<b>43.6</b> $\pm$ 1.2	<b>32.9</b> $\pm$ 1.1

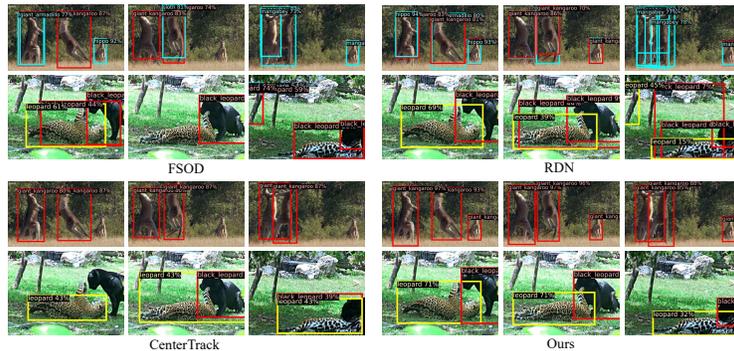
by the more powerful multi-relation MN [31] can significantly improve the performance. When cooperating with TPN, our proposed TMN outperforms MN by 3.2 AP in the temporal domain using aligned query features. The improved TMN+ reaches 30.0 AP performance by capitalizing on better generalization and representative feature, which is optimized with the label-smoothing regularization and support classification module, bringing about respectively 1.5 and 3.0 performance increase. Note that our support classification module is fundamentally different from the meta-loss in Meta R-CNN [129] which requires training on novel classes to avoid prediction ambiguity in object attentive vectors, while our method targets at generating more representative features in the Euclidean space to generalize better on novel classes without any fine-tuning.

### 5.3 Advantages of Temporal Matching

Temporal matching has two substantial advantages over image-based matching: **Ghost Proposal Removal.** Image-based matching suffers heavily from “ghost proposals” which are hard background proposals with similar appearance to foreground proposals. It is difficult to filter them out by the RPN in the spatial domain due to appearance ambiguity, while much easier to distinguish in the temporal domain due to their intermittent “ghost” or discontinuous appearances across frames. Our TPN takes this advantage to get rid of ghost proposals and thus obtains better detection performance.

**Representative Feature.** From the feature perspective, image-based matching exploits proposal features from each query frame to match with supports individually. Such independent query feature is inadequate in representing a target video object, especially those in bad visual quality due to *e.g.*, large deformation, motion blur or heavy occlusion, thus is liable to bad comparison results in the subsequent matching procedure and leading to bad predictions. In contrast, our temporal matching aggregates object features across frames in the tube proposal into a robust representative feature for the target video object, which helps the subsequent matching procedure to produce better result.

**Validation.** We show quantitatively and qualitatively the above advantages of our temporal matching. Specifically, we transform our tube-based matching to the image-based matching by performing per-frame detection and matching during inference. With the same trained model, the performance drastically drops



**Fig. 4. Qualitative 5-shot detection results on novel classes of FSVOD dataset.** Our tube-based approach successfully detects objects in novel classes, while other methods miss or misclassify target objects or detect ghost objects.

from 30.0 to 25.8 after replacing tube-based feature by image-based feature. The large performance gap indicates the effectiveness of tube-based matching in the FSVOD task. In Figure 4, the image-based methods produce ghost proposals and fails the target object matching, while our approach produces much better performance without suffering from ghost proposals.

#### 5.4 Object Indexing in Massive Videos

Our FSVOD task enables models properly solving the object indexing/retrieval problem in massive videos, which is infeasible or extreme hard for other computer vision tasks. Specifically, we retrieve video clips for the target support class if there exists a detected box with the class score larger than 0.05. Thanks to the full-way evaluation, our FSVOD actually performs indexing for every class in the entire video set. We use the widely-used  $F_1$  score to evaluate the retrieval performance. Our FSVOD model achieves 0.414  $F_1$  score on FSVOD-500 test set, while the classic few-shot object detection model [31] only obtains 0.339  $F_1$  score because of its numerous false positive predictions in videos. More details are in the supplementary material.

## 6 Conclusion

This paper proposes FSVOD for detecting objects in novel classes in a query *video* given *only a few support images*. FSVOD can be applied in high diversity/dynamic scenarios for solving relevant real-world problem that is infeasible or hard for other computer vision tasks. We contribute a new large-scale, class-balanced FSVOD dataset, which contains 500 classes of objects in high diversity with high-quality annotations. Our tube proposal network and aligned matching network effectively employ the temporal information in proposal generation and matching. Extensive comparison have been performed to compare related methods on two datasets to validate that our FSVOD method produces the best performance. We hope this paper will kindle future FSVOD research.

## References

1. Bell, S., Lawrence Zitnick, C., Bala, K., Girshick, R.: Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In: CVPR (2016) [4](#)
2. Bergmann, P., Meinhardt, T., Leal-Taixe, L.: Tracking without bells and whistles. In: ICCV (2019) [4](#)
3. Bertasius, G., Torresani, L., Shi, J.: Object detection in video with spatiotemporal sampling networks. In: ECCV (2018) [1](#), [4](#)
4. Bertinetto, L., Henriques, J.F., Torr, P.H., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. In: ICLR (2019) [3](#)
5. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: ECCV (2016) [1](#), [4](#)
6. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: ICCV (2019) [4](#)
7. Bhat, G., Danelljan, M., Van Gool, L., Timofte, R.: Know your surroundings: Exploiting scene information for object tracking. In: ECCV (2020) [4](#)
8. Cai, Z., Fan, Q., Feris, R.S., Vasconcelos, N.: A unified multi-scale deep convolutional neural network for fast object detection. In: ECCV (2016) [4](#)
9. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: CVPR (2018) [4](#)
10. Cao, K., Ji, J., Cao, Z., Chang, C.Y., Niebles, J.C.: Few-shot video classification via temporal alignment. In: CVPR (2020) [2](#)
11. Chen, H., Wang, Y., Wang, G., Qiao, Y.: Lstd: A low-shot transfer detector for object detection. In: AAAI (2018) [3](#)
12. Chen, Y., Cao, Y., Hu, H., Wang, L.: Memory enhanced global-local aggregation for video object detection. In: CVPR (2020) [4](#), [11](#), [12](#)
13. Chu, P., Ling, H.: Famnet: Joint learning of feature, affinity and multi-dimensional assignment for online multiple object tracking. In: ICCV (2019) [1](#), [4](#)
14. Dai, J., Li, Y., He, K., Sun, J.: R-fcn: Object detection via region-based fully convolutional networks. In: NeurIPS (2016) [4](#)
15. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: ICCV (2017) [6](#)
16. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: ATOM: accurate tracking by overlap maximization. In: CVPR (2019) [4](#)
17. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: CVPR (2017) [4](#)
18. Danelljan, M., Gool, L.V., Timofte, R.: Probabilistic regression for visual tracking. In: CVPR (2020) [4](#)
19. Danelljan, M., Hager, G., Shahbaz Khan, F., Felsberg, M.: Learning spatially regularized correlation filters for visual tracking. In: ICCV (2015) [4](#)
20. Danelljan, M., Shahbaz Khan, F., Felsberg, M., Van de Weijer, J.: Adaptive color attributes for real-time visual tracking. In: CVPR (2014) [4](#)
21. Dave, A., Khurana, T., Tokmakov, P., Schmid, C., Ramanan, D.: Tao: A large-scale benchmark for tracking any object. In: ECCV (2020) [9](#)
22. Deng, H., Hua, Y., Song, T., Zhang, Z., Xue, Z., Ma, R., Robertson, N., Guan, H.: Object guided external memory network for video object detection. In: ICCV (2019) [4](#)
23. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009) [9](#), [11](#)

24. Deng, J., Pan, Y., Yao, T., Zhou, W., Li, H., Mei, T.: Relation distillation networks for video object detection. In: ICCV (2019) 4, 11, 12
25. Dollár, P., Wojek, C., Schiele, B., Perona, P.: Pedestrian detection: A benchmark. In: CVPR (2009) 10
26. Dong, N., Xing, E.P.: Few-shot semantic segmentation with prototype learning. In: BMVC (2018) 3
27. Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., Tian, Q.: Centernet: Keypoint triplets for object detection. In: ICCV (2019) 4
28. Ess, A., Leibe, B., Schindler, K., Van Gool, L.: A mobile vision system for robust multi-person tracking. In: CVPR (2008) 10
29. Fan, H., Lin, L., Yang, F., Chu, P., Deng, G., Yu, S., Bai, H., Xu, Y., Liao, C., Ling, H.: Lasot: A high-quality benchmark for large-scale single object tracking. In: CVPR (2019) 4, 9
30. Fan, Q., Ke, L., Pei, W., Tang, C.K., Tai, Y.W.: Commonality-parsing network across shape and appearance for partially supervised instance segmentation. In: ECCV (2020) 3
31. Fan, Q., Zhuo, W., Tang, C.K., Tai, Y.W.: Few-shot object detection with attention-rpn and multi-relation detector. In: CVPR (2020) 2, 3, 5, 7, 11, 12, 13, 14
32. Fan, Z., Ma, Y., Li, Z., Sun, J.: Generalized few-shot object detection without forgetting. In: CVPR (2021) 3
33. Fang, K., Xiang, Y., Li, X., Savarese, S.: Recurrent autoregressive networks for online multi-object tracking. In: WACV (2018) 4
34. Feichtenhofer, C., Pinz, A., Zisserman, A.: Detect to track and track to detect. In: ICCV (2017) 4
35. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML (2017) 3, 5
36. Girshick, R.: Fast r-cnn. In: ICCV (2015) 1, 3
37. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: CVPR (2014) 4
38. Goutam, B., Järemo Lawin, F., Danelljan, M., Robinson, A., Felsberg, M., Van Gool, L., Timofte, R.: Learning what to learn for video object segmentation. In: ECCV (2020) 4
39. Gui, L.Y., Wang, Y.X., Ramanan, D., Moura, J.M.F.: Few-shot human motion prediction via meta-learning. In: ECCV (2018) 3
40. Guo, Q., Feng, W., Zhou, C., Huang, R., Wan, L., Wang, S.: Learning dynamic siamese network for visual object tracking. In: ICCV (2017) 4
41. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: CVPR (2019) 9
42. Hariharan, B., Girshick, R.: Low-shot visual recognition by shrinking and hallucinating features. In: ICCV (2017) 3
43. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017) 4
44. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) 11
45. Held, D., Thrun, S., Savarese, S.: Learning to track at 100 fps with deep regression networks. In: ECCV (2016) 4
46. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. IEEE TPAMI (2014) 4
47. Hu, H., Bai, S., Li, A., Cui, J., Wang, L.: Dense relation distillation with context-aware aggregation for few-shot object detection. In: CVPR (2021) 3

48. Hu, T., Pengwan, Zhang, C., Yu, G., Mu, Y., Snoek, C.G.M.: Attention-based multi-context guiding for few-shot semantic segmentation. In: AAAI (2019) [3](#)
49. Huang, L., Zhao, X., Huang, K.: Got-10k: A large high-diversity benchmark for generic object tracking in the wild. IEEE TPAMI (2019) [9](#)
50. Kang, B., Liu, Z., Wang, X., Yu, F., Feng, J., Darrell, T.: Few-shot object detection via feature reweighting. In: ICCV (2019) [3](#)
51. Kang, K., Li, H., Xiao, T., Ouyang, W., Yan, J., Liu, X., Wang, X.: Object detection in videos with tubelet proposal networks. In: CVPR (2017) [4](#)
52. Kang, K., Ouyang, W., Li, H., Wang, X.: Object detection from video tubelets with convolutional neural networks. In: CVPR (2016) [4](#)
53. Karlinsky, L., Shtok, J., Harary, S., Schwartz, E., Aides, A., Feris, R., Giryes, R., Bronstein, A.M.: Repmet: Representative-based metric learning for classification and few-shot object detection. In: CVPR (2019) [3](#)
54. Kemelmacher-Shlizerman, I., Seitz, S.M., Miller, D., Brossard, E.: The megaface benchmark: 1 million faces for recognition at scale. In: CVPR (2016) [10](#)
55. Khodadadeh, S., Boloni, L., Shah, M.: Unsupervised meta-learning for few-shot image classification. NeurIPS (2019) [2](#)
56. Kim, C., Li, F., Ciptadi, A., Rehg, J.M.: Multiple hypothesis tracking revisited. In: ICCV (2015) [4](#)
57. Kim, D., Woo, S., Lee, J.Y., Kweon, I.S.: Video panoptic segmentation. In: CVPR (2020) [4](#)
58. Koch, G., Zemel, R., Salakhutdinov, R.: Siamese neural networks for one-shot image recognition. In: ICMLW (2015) [2, 3, 5](#)
59. Kong, T., Sun, F., Liu, H., Jiang, Y., Li, L., Shi, J.: Foveabox: Beyond anchor-based object detection. IEEE TIP (2020) [4](#)
60. Kong, T., Sun, F., Yao, A., Liu, H., Lu, M., Chen, Y.: Ron: Reverse connection with objectness prior networks for object detection. In: CVPR (2017) [4](#)
61. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Cehovin Zajc, L., Vojir, T., Hager, G., Lukezic, A., Eldesokey, A., et al.: The visual object tracking vot2017 challenge results. In: ICCVW (2017) [4](#)
62. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Cehovin Zajc, L., Vojir, T., Bhat, G., Lukezic, A., Eldesokey, A., et al.: The sixth visual object tracking vot2018 challenge results. In: ECCV (2018) [4](#)
63. Kristan, M., Matas, J., Leonardis, A., Felsberg, M., Cehovin, L., Fernandez, G., Vojir, T., Hager, G., Nebehay, G., Pflugfelder, R.: The visual object tracking vot2015 challenge results. In: ICCVW (2015) [4](#)
64. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. In: ECCV (2018) [4](#)
65. Lee, Y., Choi, S.: Gradient-based meta-learning with learned layerwise metric and subspace. In: ICML (2018) [3](#)
66. Li, A., Li, Z.: Transformation invariant few-shot object detection. In: CVPR (2021) [3](#)
67. Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: Siamrpn++: Evolution of siamese visual tracking with very deep networks. In: CVPR (2019) [1, 4](#)
68. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: CVPR (2018) [4](#)
69. Li, B., Yang, B., Liu, C., Liu, F., Ji, R., Ye, Q.: Beyond max-margin: Class margin equilibrium for few-shot object detection. In: CVPR (2021) [3](#)
70. Li, X., Wei, T., Chen, Y.P., Tai, Y.W., Tang, C.K.: Fss-1000: A 1000-class dataset for few-shot segmentation. In: CVPR (2020) [3, 5](#)

71. Li, Y., Chen, Y., Wang, N., Zhang, Z.: Scale-aware trident networks for object detection. In: ICCV (2019) 4
72. Li, Y., Zhu, H., Cheng, Y., Wang, W., Teo, C.S., Xiang, C., Vadakkepat, P., Lee, T.H.: Few-shot object detection via classification refinement and distractor retreatment. In: CVPR (2021) 3
73. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017) 3
74. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017) 4
75. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014) 10, 11
76. Liu, S., Huang, D., et al.: Receptive field block net for accurate and fast object detection. In: ECCV (2018) 4
77. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: ECCV (2016) 4
78. Liu, W., Liao, S., Ren, W., Hu, W., Yu, Y.: High-level semantic feature detection: A new perspective for pedestrian detection. In: CVPR (2019) 4
79. Liu, Y., Zhang, X., Zhang, S., He, X.: Part-aware prototype network for few-shot semantic segmentation. In: ECCV (2020) 3
80. Lu, X., Li, B., Yue, Y., Li, Q., Yan, J.: Grid r-cnn. In: CVPR (2019) 4
81. Lu, Z., Rathod, V., Votel, R., Huang, J.: Retinatrack: Online single stage joint detection and tracking. In: CVPR (2020) 4
82. Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixé, L., Leibe, B.: Hota: A higher order metric for evaluating multi-object tracking. IJCV (2021) 4
83. Mayer, C., Danelljan, M., Paudel, D.P., Gool, L.V.: Learning target candidate association to keep track of what not to track. In: ICCV (2021) 4
84. Michaelis, C., Bethge, M., Ecker, A.S.: One-shot segmentation in clutter. In: ICML (2018) 3
85. Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: Mot16: A benchmark for multi-object tracking. arXiv preprint arXiv:1603.00831 (2016) 9
86. Muller, M., Bibi, A., Giancola, S., Alsubaihi, S., Ghanem, B.: Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In: ECCV (2018) 9
87. Müller, R., Kornblith, S., Hinton, G.E.: When does label smoothing help? In: NeurIPS (2019) 8
88. Oh, S.W., Lee, J.Y., Xu, N., Kim, S.J.: Video object segmentation using space-time memory networks. In: ICCV (2019) 1, 4
89. Pang, B., Li, Y., Zhang, Y., Li, M., Lu, C.: Tubetk: Adopting tubes to track multi-object in a one-step training model. In: CVPR (2020) 4
90. Peng, J., Wang, C., Wan, F., Wu, Y., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., Fu, Y.: Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking. In: ECCV (2020) 1, 4, 7, 11, 12
91. Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: CVPR (2016) 4
92. Perez-Rua, J.M., Zhu, X., Hospedales, T.M., Xiang, T.: Incremental few-shot object detection. In: CVPR (2020) 3

93. Real, E., Shlens, J., Mazzocchi, S., Pan, X., Vanhoucke, V.: Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In: CVPR (2017) 9
94. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: CVPR (2016) 4
95. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: CVPR (2017) 4
96. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NeurIPS (2015) 1, 3, 6, 11, 12
97. Sadeghian, A., Alahi, A., Savarese, S.: Tracking the untrackable: Learning to track multiple cues with long-term dependencies. In: ICCV (2017) 4
98. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T.: Meta-learning with memory-augmented neural networks. In: ICML (2016) 5
99. Sbai, O., Couprie, C., Aubry, M.: Impact of base dataset design on few-shot image classification. In: ECCV (2020) 5
100. Shao, D., Xiong, Y., Zhao, Y., Huang, Q., Qiao, Y., Lin, D.: Find and focus: Retrieve and localize video events with natural language queries. In: ECCV (2018) 4
101. Shen, Z., Liu, Z., Li, J., Jiang, Y.G., Chen, Y., Xue, X.: Dsod: Learning deeply supervised object detectors from scratch. In: ICCV (2017) 4
102. Shvets, M., Liu, W., Berg, A.C.: Leveraging long-range temporal relationships between proposals for video object detection. In: ICCV (2019) 4
103. Singh, B., Najibi, M., Davis, L.S.: Sniper: Efficient multi-scale training. In: NeurIPS (2018) 4
104. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: NeurIPS (2017) 3
105. Sun, B., Li, B., Cai, S., Yuan, Y., Zhang, C.: Fsce: Few-shot object detection via contrastive proposal encoding. In: CVPR (2021) 3
106. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016) 8
107. Tang, P., Wang, C., Wang, X., Liu, W., Zeng, W., Wang, J.: Object detection in videos by high quality object linking. IEEE TPAMI (2019) 4
108. Tang, P., Wang, X., Bai, X., Liu, W.: Multiple instance detection network with online instance classifier refinement. In: CVPR (2017) 4
109. Tao, R., Gavves, E., Smeulders, A.W.: Siamese instance search for tracking. In: CVPR (2016) 4
110. Thulasidasan, S., Chennupati, G., Bilmes, J.A., Bhattacharya, T., Michalak, S.: On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In: NeurIPS (2019) 8
111. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: ICCV (2019) 4
112. Valmadre, J., Bertinetto, L., Henriques, J., Vedaldi, A., Torr, P.H.: End-to-end representation learning for correlation filter based tracking. In: CVPR (2017) 4
113. Valmadre, J., Bewley, A., Huang, J., Sun, C., Sminchisescu, C., Schmid, C.: Local metrics for multi-object tracking. arXiv preprint arXiv:2104.02631 (2021) 4
114. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: NeurIPS (2016) 2, 3
115. Wang, S., Zhou, Y., Yan, J., Deng, Z.: Fully motion-aware network for video object detection. In: ECCV (2018) 4
116. Wang, X., Huang, T.E., Darrell, T., Gonzalez, J.E., Yu, F.: Frustratingly simple few-shot object detection. In: ICML (2020) 3

117. Wang, Y.X., Girshick, R., Hebert, M., Hariharan, B.: Low-shot learning from imaginary data. In: CVPR (2018) [3](#)
118. Wang, Y.X., Ramanan, D., Hebert, M.: Meta-learning to detect rare objects. In: CVPR (2019) [3](#)
119. Woo, S., Kim, D., Cho, D., Kweon, I.S.: Linknet: Relational embedding for scene graph. In: NeurIPS (2018) [4](#)
120. Wu, C.Y., Feichtenhofer, C., Fan, H., He, K., Krahenbuhl, P., Girshick, R.: Long-term feature banks for detailed video understanding. In: CVPR (2019) [4](#)
121. Wu, H., Chen, Y., Wang, N., Zhang, Z.: Sequence level semantics aggregation for video object detection. In: ICCV (2019) [4](#)
122. Wu, J., Liu, S., Huang, D., Wang, Y.: Multi-scale positive sample refinement for few-shot object detection. In: ECCV (2020) [3](#)
123. Wu, Y., Lim, J., Yang, M.H.: Online object tracking: A benchmark. In: CVPR (2013) [4](#)
124. Xiao, F., Jae Lee, Y.: Video object detection with an aligned spatial-temporal memory. In: ECCV (2018) [4](#)
125. Xiao, T., Li, S., Wang, B., Lin, L., Wang, X.: Joint detection and identification feature learning for person search. In: CVPR (2017) [10](#)
126. Xiao, Y., Marlet, R.: Few-shot object detection and viewpoint estimation for objects in the wild. In: ECCV (2020) [3](#)
127. Xu, J., Cao, Y., Zhang, Z., Hu, H.: Spatial-temporal relation networks for multi-object tracking. In: ICCV (2019) [4](#)
128. Xu, N., Yang, L., Fan, Y., Yang, J., Yue, D., Liang, Y., Price, B., Cohen, S., Huang, T.: Youtube-vos: Sequence-to-sequence video object segmentation. In: ECCV (2018) [4](#), [9](#)
129. Yan, X., Chen, Z., Xu, A., Wang, X., Liang, X., Lin, L.: Meta r-cnn : Towards general solver for instance-level low-shot learning. In: ICCV (2019) [3](#), [13](#)
130. Yang, F.S.Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR (2018) [3](#), [12](#)
131. Yang, L., Fan, Y., Xu, N.: Video instance segmentation. In: ICCV (2019) [4](#), [9](#), [11](#)
132. Yang, Y., Wei, F., Shi, M., Li, G.: Restoring negative information in few-shot object detection. In: NeurIPS (2020) [3](#)
133. Yang, Z., Liu, S., Hu, H., Wang, L., Lin, S.: Reppoints: Point set representation for object detection. In: ICCV (2019) [4](#)
134. Yang, Z., Wang, Y., Chen, X., Liu, J., Qiao, Y.: Context-transformer: tackling object confusion for few-shot detection. In: AAAI (2020) [2](#), [3](#)
135. Yu, F., Li, W., Li, Q., Liu, Y., Shi, X., Yan, J.: Poi: Multiple object tracking with high performance detection and appearance feature. In: ECCV (2016) [4](#)
136. Zhan, Y., Wang, C., Wang, X., Zeng, W., Liu, W.: A simple baseline for multi-object tracking. IJCV (2021) [4](#), [11](#), [12](#)
137. Zhang, C., Cai, Y., Lin, G., Shen, C.: Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In: CVPR (2020) [3](#)
138. Zhang, L., Zhou, S., Guan, J., Zhang, J.: Accurate few-shot object detection with support-query mutual guidance and hybrid loss. In: CVPR (2021) [3](#)
139. Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In: CVPR (2020) [4](#)
140. Zhang, W., Wang, Y.X.: Hallucination improves few-shot object detection. In: CVPR (2021) [3](#)

141. Zhang, Z., Qiao, S., Xie, C., Shen, W., Wang, B., Yuille, A.L.: Single-shot object detection with enriched semantics. In: CVPR (2018) [4](#)
142. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. In: ECCV (2020) [4](#), [11](#), [12](#)
143. Zhou, X., Zhuo, J., Krahenbuhl, P.: Bottom-up object detection by grouping extreme and center points. In: CVPR (2019) [4](#)
144. Zhu, C., Chen, F., Ahmed, U., Savvides, M.: Semantic relation reasoning for shot-stable few-shot object detection. In: CVPR (2021) [3](#)
145. Zhu, J., Yang, H., Liu, N., Kim, M., Zhang, W., Yang, M.H.: Online multi-object tracking with dual matching attention networks. In: ECCV (2018) [4](#)
146. Zhu, L., Yang, Y.: Compound memory networks for few-shot video classification. In: ECCV (2018) [2](#)
147. Zhu, R., Zhang, S., Wang, X., Wen, L., Shi, H., Bo, L., Mei, T.: Scratchdet: Training single-shot object detectors from scratch. In: CVPR (2019) [4](#)
148. Zhu, X., Dai, J., Yuan, L., Wei, Y.: Towards high performance video object detection. In: CVPR (2018) [4](#)
149. Zhu, X., Wang, Y., Dai, J., Yuan, L., Wei, Y.: Flow-guided feature aggregation for video object detection. In: ICCV (2017) [4](#)
150. Zhu, X., Xiong, Y., Dai, J., Yuan, L., Wei, Y.: Deep feature flow for video recognition. In: CVPR (2017) [4](#)