

UnOVOST: Unsupervised Offline Video Object Segmentation and Tracking

Jonathon Luiten* Idil Esen Zulfikar* Bastian Leibe
Computer Vision Group, RWTH Aachen University

{luiten,leibe}@vision.rwth-aachen.de idil.esen.zuelfikar@rwth-aachen.de

Abstract

We address *Unsupervised Video Object Segmentation (UVOS)*, the task of automatically generating accurate pixel masks for salient objects in a video sequence and of tracking these objects consistently through time, without any input about which objects should be tracked. Towards solving this task, we present *UnOVOST (Unsupervised Offline Video Object Segmentation and Tracking)* as a simple and generic algorithm which is able to track and segment a large variety of objects. This algorithm builds up tracks in a number of stages, first grouping segments into short tracklets that are spatio-temporally consistent, before merging these tracklets into long-term consistent object tracks based on their visual similarity. In order to achieve this we introduce a novel tracklet-based Forest Path Cutting data association algorithm which builds up a decision forest of track hypotheses before cutting this forest into paths that form long-term consistent object tracks. When evaluating our approach on the DAVIS 2017 Unsupervised dataset we obtain state-of-the-art performance with a mean $\mathcal{J}\&\mathcal{F}$ score of 67.9% on the val, 58% on the test-dev and 56.4% on the test-challenge benchmarks, obtaining first place in the DAVIS 2019 Unsupervised Video Object Segmentation Challenge. *UnOVOST* even performs competitively with many semi-supervised video object segmentation algorithms even though it is not given any input as to which objects should be tracked and segmented.

1. Introduction

Video Object Segmentation (VOS) aims at automatically generating accurate pixel masks for objects in each frame of a video, then associating those proposed object pixel masks in the successive frames to obtain temporally consistent tracks. VOS has mostly been tackled in a semi-supervised fashion [19, 32, 30], where the object masks of the objects to be tracked in the first-frame are given, and only those objects need to be tracked and segmented throughout the rest of the video.

*Equal Contribution



Figure 1. Example results of UnOVOST on three sequences from the DAVIS Unsupervised Dataset. UnOVOST is able to accurately segment and track many diverse objects simultaneously. Frames $1, \frac{T}{4}, \frac{T}{2}, \frac{3T}{4}$ and T are shown, where T is the number of frames in the sequence.

In this paper we tackle VOS in the more general unsupervised setting [5]. In such a setting we need to detect all of the possible objects in the video and track and segment them throughout the whole video. Results of our method on this task can be seen in Figure 1. In this setting, methods are evaluating against a possibly incomplete set of ground-truth objects. As such methods are not penalized for segmenting more objects than present in the ground-truth. However, the number of predictions that can be made is limited in that predicted masks may not have overlapping pixels, and a maximum number of objects may be proposed across a whole video. As such UVOS methods must seek to segment and track the most salient objects in a video regardless of the category of those objects. Saliency here is defined as the objects that catch and maintain the gaze of a viewer across the whole of the video sequence. The definition of an object is also important and possibly ambiguous. Two important factors in determining objectness are that objects should consistently have common fate, moving together consistently throughout the scene, and that they should also be semantically consistent.

An algorithm that tackles the UVOS task has many interesting real-world applications. One such example is in robotics and autonomous vehicles where it is of crucial importance to be able to understand the precise location and motion of a huge variety of objects, from far more categories than present in any labeled dataset.

To solve this UVOS task, we present the UnOVOST (Unsupervised Offline Video Object Segmentation and Tracking) algorithm. This algorithm hierarchically builds up object segmentation tracks in multiple stages (see Figure 2). After obtaining a set of candidate object proposal masks per frame using Mask R-CNN [9]. We then reduce the set of mask proposals to a set which does not contain overlapping pixels by sub-selecting and clipping the given proposals. In order to perform segment tracking we use two main similarity cues, the spatio-temporal consistency of the mask segments in contiguous frames under optical flow warping, and the appearance-based visual similarity of the mask segments encoded as an object re-identification vector.

We then develop a novel data-association algorithm using these two similarity cues which accurately merges these mask segments into tracks. Our algorithm works in two stages. The first stage uses only the spatio-temporal consistency cues to merge segments in contiguous frames into short-tracklets which contain segments that are very likely to belong to the same object. In a second stage, we then merge these short-tracklets into long-term consistent tracks using their visual similarity. This two stage process has the benefit that the easy tracking decisions are made early and then fixed, reducing the size of the required search-space for performing data-association, and enabling information to be pooled over segments within a tracklet to better model object properties used for tracking.

For the second-stage, we propose a novel Forest Path Cutting (FPC) algorithm. This algorithm builds a forest consisting of decision trees of possible track hypotheses. The final set of object tracks is then produced by iteratively cutting paths from this forest, until the forest is divided into a non-conflicting set of paths which are the final tracks. This algorithm is both simple and efficient while being powerful enough to model the combinatorial complexity of the long-term data-association problem.

When evaluating UnOVOST on the unsupervised DAVIS benchmark dataset [5] we achieve state-of-the-art results compared to all previous methods, as well as results competitive with semi-supervised methods using the given first-frame mask as guidance for which objects to track and segment. Our method also achieves the first place in the DAVIS 2019 Unsupervised Video Object Segmentation Challenge. When extending our method to the task of Video Instance Segmentation (VIS) by adding classifying our object tracks, we also obtain state-of-the-art results on the YouTube-VIS benchmark and won the 2019 YouTube-VIS challenge.

2. Related Work

Multi-Object Unsupervised Video Object Segmentation. The UVOS task (also known as zero-shot VOS) is quite recent, and there are few methods that tackle this task. UVOS [5] was proposed as a challenge task for the 2019 DAVIS Challenge on Video Object Segmentation. [5] evaluate the RVOS (Recurrent Video Object Segmentation) [29] method for the UVOS task. This method uses a number of recurrent neural networks, one over the set of objects, and one over time to generate tracks. Our method, UnOVOST, outperforms RVOS by more than 25 percentage points on the $\mathcal{J}\&\mathcal{F}$ metric on all benchmarks. In the 2019 DAVIS Challenge, our method obtained first place. The second [41] and third [7] place methods presented very different approaches to the UVOS task. [41] propose to run a detector on each frame, as well as a series of single object trackers used to merge the detections into tracks. [7] adapts [19] from semi-supervised VOS to UVOS task, while adding a proposal pruning step after a number of initial frames, and then tracking these objects as though this was a semi-supervised task.

Single-Object Unsupervised Video Object Segmentation. There has been a number of papers tackling single-object unsupervised video object segmentation (SOUVOS) [15, 16, 12, 27, 28]. This is inherently a different problem to the multi-object task that we tackle in this paper. SOUVOS is closer related to foreground/background segmentation as it requires only one foreground area to be segmented which often is a grouping of multiple objects into one foreground object. This task is often evaluated on the DAVIS 2016 single object benchmark [24]. This task requires estimating the single most salient grouping of foreground objects in a video. Methods that tackle this task, such as [12] and [28] often perform two class segmentation on an image concatenated with optical-flow.

Motion Segmentation. Another related field is motion segmentation. This task differs from video object segmentation in that it only requires the segmentation of objects that are moving [2], whereas UVOS requires the segmentation of all objects whether they are moving or not. Motion segmentation methods [3, 37, 8] are often based on low-level vision features such as the optical-flow. [8] adapts Mask R-CNN [9] to operate on both image and optical-flow input. [37] extract features from the combination of the image and the optical-flow and clusters these. [3] develops a two-stage model that estimates piece-wise rigid motions, which are then merged into objects. This is evaluated as either a multi-object task, or a foreground/background estimation task often using the FSMB [21] dataset.

Multi-Object Semi-Supervised Video Object Segmentation. Semi-Supervised Video Object Segmentation (SSVOS) is where the objects that need to be tracked are given as segmentation masks in the first frame. Algorithms

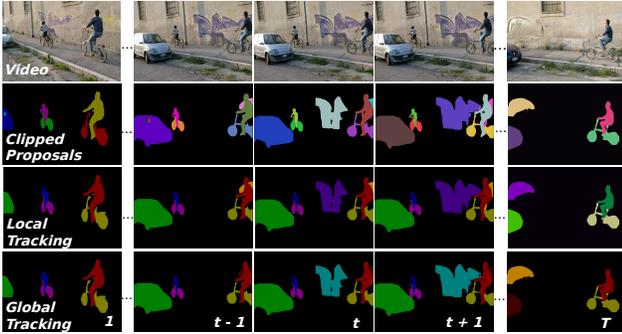


Figure 2. An overview of the UnOVOST algorithm. From an input video (row 1) a number of object mask proposals per frame are generated, sub-selected and clipped to have non-overlapping pixels (row 2). These are grouped into short-tracklets using the spatio-temporal consistency of these segments under optical-flow (row 3). These tracklets are then merged into long-term consistent object tracks using the tracklets’ visual similarity and our novel Forest Path Cutting (FPC) data association algorithm.

that tackle this task often finetune a segmentation network on the given first frame [4, 32, 19], or propagate from the given first-frame mask directly to the rest of the video [30, 36]. These methods are not able to be easily adapted to UVOS as they rely heavily on the first-frame mask. [19] is the closest related SSVOS to our method, as it also produces generic object segmentation proposals and links these in time with spatio-temporal and visual similarity cues. However, unlike our method, [19] finetunes all of its components heavily on the first-frame, uses the given-first frame to guide which objects to track, and performs data association in a simple frame-by-frame fashion. SSVOS is often evaluated on the DAVIS 2017 semi-supervised dataset.

Video Instance Segmentation and Multi-Object Tracking and Segmentation. Recently, the related tasks of Video Instance Segmentation (VIS) [39] and Multi-Object Tracking and Segmentation (MOTS) [31] has been proposed. These tasks are similar to UVOS in that objects need to be tracked and segmented without being given guidance on which particular instances are to be tracked. However, these tasks differs from UVOS in that only objects belonging to a specified categories need to be tracked and segmented, as well as these being classified correctly. This significantly simplifies the task, and limits the applicability of methods that tackle these tasks. MOTS differs from VIS in that in MOTS sequences are much longer and many more instances are present with objects disappearing and re-appearing much more often. MOTS is evaluated on the KITTI and MOTChallenge datasets [31]. VIS on the YouTubeVIS benchmark [39]. We extend UnOVOST from the UVOS task to the VIS task by classifying our resulting tracks, and also achieve state-of-the-art performance on this task.

Category Agnostic Multi-Object Tracking. Previously, a number of methods [22, 23] have attempted to extent multi-

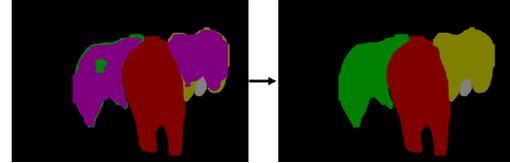


Figure 3. Initial stage of the UnOVOST algorithm. Mask proposals are generated by Mask R-CNN with a low scoring threshold that generates a large number of overlapping proposals. These are then sub-selected and clipped based on their score and intersection to produce a set of non-overlapping mask proposals in each frame.

object tracking methods beyond tracking objects from a pre-defined set of object categories. These methods [22, 23] have typically relied on the presence of stereo-camera input to obtain 3D information for evaluating the objectness of generic object proposals. These methods also make no attempt to create a set of object tracks without overlapping segment masks, instead they create a large set of track proposals that have large overlap with one-another, often tracking the same object multiple times on different scales. Our method by contrast works on monocular video and creates a set of segmentation tracks without overlapping segment masks.

Data-Association for Multi-Object Tracking. The task of multi-object tracking (MOT) has a long research history [14]. The leading paradigm for MOT has become tracking-by-detection, where a set of object detections are proposed, and tracking is reduced to a data-association problem. In this paper we propose a new data-association algorithm designed specifically for the UVOS task. Previous data-association methods are either too simple or unnecessarily complex. Many methods such as those used in [31] and [19] only take into account associations in the previous frame, or the previous and first frame [19], and do not use the context from the whole video. On the other hand, data-association algorithms such as [13] are unnecessarily complex in that they produce an exponentially large number of potential track hypotheses and score each of these individually. Our data-association algorithm is able to take advantage of a number of simplifications to be able to use the whole video context to evaluate the likelihood of a tracking hypothesis, while being much simpler and efficient. We take advantage of the fact that mask segments can not overlap to significantly reduce the set of possible tracking combinations. Furthermore, we split tracking into two components, firstly grouping proposals based on spatio-temporal consistency, before only using visual similarity in a second stage. These simplifications in combination with our efficient Forest Path Cutting algorithm, results in an algorithm that is accurate, powerful and efficient.

3. Approach

In this section we detail the specifics of our novel UnOVOST algorithm for tackling the UVOS task. Our method begins by generating a large set of generic object proposals and sub-selects and clips these to be non-overlapping. These proposals are then grouped over contiguous frames into short-tracklets based on the spatio-temporal consistency of the object proposals under an optical-flow warping. We then merge these tracklets into long-term consistent object tracks using our novel Forest Path Cutting algorithm and visual similarity cues between tracklets. Finally, a final set of object tracks is selected based on their video saliency.

Unlike previous approaches [19, 31, 8] our algorithm is able to segment and track objects regardless of their object class, whether they are static or undergo motion, and whether they are present in the foreground or background of the scene. Our algorithm instead relies on a more general concept of objectness to determine what should consist of an object to be tracked and segmented. An overview can be seen in Figure 2.

Object Mask Proposal Generation. We generate a large number of proposals, segmentation masks which cover potential objects, for a diverse range of objects. Specifically we use a Mask R-CNN [9] implementation by [34] with a ResNet101 [10] backbone trained on COCO [18]. Although this network has been trained to detect the 80 COCO categories, we find that when using a low-confidence threshold this network produces adequate mask proposals for objects beyond these 80 categories. This network produces masks, bounding boxes, object categories and confidence scores for object proposals as outputs. We discard the object categories and treat all detections as if they come from the same foreground object category. We extract all proposal masks with a confidence score greater than 0.1. An example of generated overlapping proposals can be found in Figure 3.

Proposal Sub-Selection and Clipping. In order to simplify the tracking and segmentation problem, we initially ensure that our set of segment masks do not overlap in each frame individually before tracking these segments throughout the video. All proposal masks in a frame are compared against one another using their intersection over union (IoU) to detect overlaps. If the IoU between two proposal masks is higher than 0.2, then the proposal mask with higher confidence score is kept and the other proposal mask is removed. This is a form of mask-based non-maximum suppression. For all the remaining masks, we clip overlaps so that the mask with the highest confidence score is on top of a mask with a lower score. This set of proposals without overlaps has three advantages, it reduces the number of proposals that need to be tracked, simplifies the matching using temporal consistency cues as conflicting proposals are removed, and removes many spurious masks not belonging to real objects. See Figure 3.

Algorithm 1: Forest Path Cutting Algorithm (FPC)

Data: Tracklets L_i with average ReID vectors R_i , beginning timestep b_i and ending timestep e_i , ordered by increasing b_i .

Result: Tracks F_j which are groupings of tracklets.

Define: Visual Similarity $V_{i,j}$

$$V_{i,j} := 1 - \frac{\|R_i - R_j\|}{\max_{m \in \{1 \dots T\}, n \in \{1 \dots T\}} \|R_m - R_n\|}$$

Part 1: Build a forest of track hypotheses, by calculating optimal predecessors M_i for each tracklet L_i .

```

for  $i \in \{1 \dots |L|\}$  do
  if  $\{j \mid e_j < b_i\} \neq \emptyset$  then
     $k := \operatorname{argmax}_{j \mid e_j < b_i} \{V_{i,j}\}$ 
    while  $\{j \mid e_j < b_i, b_j > e_k, M_j = L_k\} \neq \emptyset$  do
       $l := \operatorname{argmax}_{j \mid e_j < b_i, j \neq k} \{V_{i,j}\}$ 
      if  $l \in \{j \mid e_j < b_i, b_j > e_k, M_j = L_k\}$  then
         $k := l$ 
      else
        break
      end
    end
     $M_i := L_k$ 
  else
     $M_i := \emptyset$ 
  end

```

end

Part 2: Define the set of track hypotheses H as the paths from root nodes to leaf nodes through the hypothesis forest, and calculate a score C_i for each path. Select final tracks F by iteratively cutting the optimal paths from the forest.

$H := \{\{L_i, M_i, M_j \mid L_j = M_i, \dots, L_k \mid M_k = \emptyset\} \mid L_i \neq M_m \forall m\}$

$F := \emptyset$

while $H \neq \emptyset$ **do**

for $H_i \in H$ **do**

$$C_i^V := \min_{m,n \mid L_m \in H_i, L_n \in H_i} \{V_{m,n}\}$$

$$C_i^T := \sum_{j \mid L_j \in H_i} e_j - b_j + 1$$

$$C_i := 0.1C_i^V + 0.9C_i^T$$

end

$k := \operatorname{argmax}_{i \mid H_i \in H} C_i$

$F := F \cup \{H_k\}$

for $H_i \in H \setminus \{H_k\}$ **do**

$H_i := H_i \setminus H_k$

end

$H := H \setminus \{H_k\}$

end

Tracklet Generation. A tracklet is a series of proposals in contiguous frames which have been merged to belong to the same object identity. In the first stage of our tracking algorithm we join proposals in contiguous frames into tracklets if they have a very high spatio-temporal consistency.

The spatio-temporal consistency score between two proposals in contiguous timesteps is calculated as the IoU between the proposal projection from the earlier frame and the proposal in the later frame. The proposal projection is the segmentation mask generated by warping a proposal by its corresponding optical flow vectors calculated using PWC-Net [26]. Effectively the projection of this proposal into the next timestep.

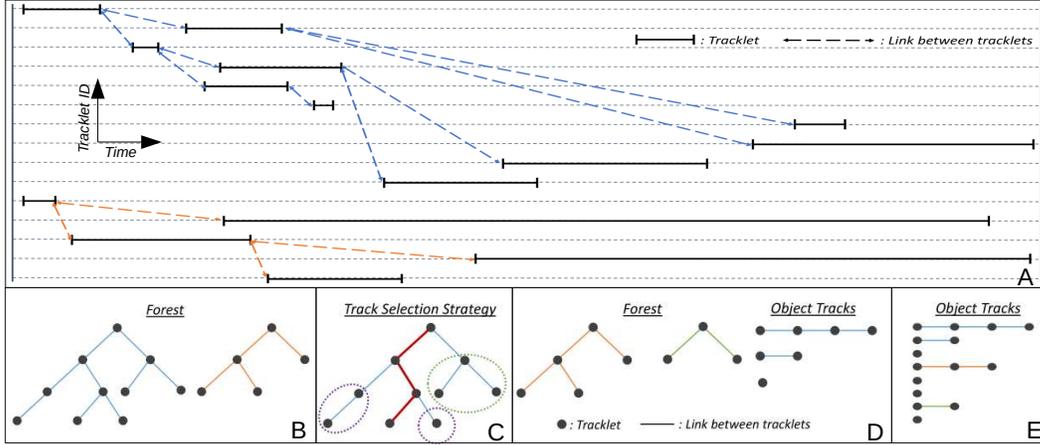


Figure 4. Visual representation of our Forest Path Cutting (FPC) algorithm. In box A, tracklets are visualized as black lines showing their temporal extent on the horizontal axis. The optimal predecessors for each tracklet are shown as blue and orange dotted lines. Box B shows that the set of tracklets with predecessors from box A defines a forest of track hypotheses. Box C shows an optimal path (in red), selected from the forest and added to the final object tracks. This optimal path is cut from its tree, dividing it into a number of sub-trees (green and purple circles). Box D shows the resulting new forest produced with sub-trees that only contain a single path added to the list of final tracks. Box E shows the final result after this process is iterated until the forest is completely divided into a set of tracks.

This first stage proceeds frame by frame. For each pair of contiguous timesteps, we create a complete bipartite graph whose nodes are the proposals in successive frames and whose edge scores are the spatio-temporal consistency scores. Edges are dropped from the graph if their score is less than 0.05. We then solve the matching problem between the two sets of nodes using the Hungarian matching algorithm, which finds an optimal set of matches between the two frames. If any proposal is not matched this ends a tracklet. Tracklets may span only a single frame.

Merging Tracklets into Tracks. A track is a set of proposals over an entire video which belong to the same object identity. A track often contains multiple tracklets with potentially frames in-between them without proposals.

The second stage of UnOVOST merges tracklets into long term tracks. To do this we introduce a novel Forest Path Cutting (FPC) algorithm. An overview of this algorithm can be seen in Algorithm 1 and Figure 4. This algorithm merges tracklets based on visual similarity cues. Note that spatio-temporal consistency now provides very little further value for data-association. If the tracklets could be easily determined to belong together by spatio-temporal consistency they would have been merged in the first stage. The remaining data association decisions are more difficult such as tracking objects through heavy or total occlusion.

For each proposal we calculate a ReID vector, a representation of the visual appearance of a proposal which can be used to compare the visual similarity of proposals or tracklets, and thus to re-identify a proposal or tracklet as belonging to a certain object identity. To calculate these vectors we use an appearance embedding network [23] which extracts an embedding from an image crop. This network is

inspired from the person re-identification community. This is a wide ResNet [35] trained with a batch-hard soft-margin version of the triplet loss. This is pretrained to distinguish classes on COCO [18], before being trained to distinguish instances on YouTube-VOS [38]. It is trained so that the embedding for instances in the same track are pulled closer together in embedding space than for difference tracks. We average ReID vectors over all proposals in a tracklet to achieve a more robust appearance representation.

To compare two tracklets we define a visual similarity score as the L2 distance of the two ReID vectors normalised to between 0 and 1, with 1 being identical, and 0 being the maximum distance between all tracklets in a video. We subtract this from one to convert it to a similarity score.

To perform long-term tracking we enumerate a set of potential track hypotheses as different combinations of tracklets. A final set of tracks can be selected as a valid subset of this set of track hypotheses.

For this task we introduce our Forest Path Cutting (FPC) algorithm as can be seen in Algorithm 1 and Figure 4. Initially (Part 1 and Box A) our algorithm builds up a forest of potential tracking hypotheses throughout the video by determining an optimal predecessor for each tracklet. Our FPC algorithm draws parallels to dynamic programming, as we wish to determine an optimal back-pointer for each tracklet to a previous predecessor tracklet. However, a naive implementation of a dynamic programming algorithm would not be able to take advantage of the desirable properties of a UVOS solution.

To calculate optimal back-pointers, our algorithm iterates over the tracklets in order from the earliest to the latest starting time. For the current tracklet L_i , if there are

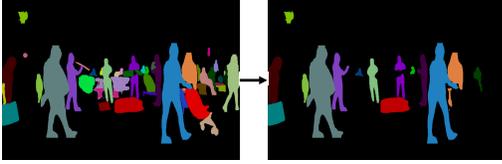


Figure 5. An example of the final stage of UnOVOST, where the final set of tracks is reduced to a maximum of 20 objects over the whole video using our video object saliency metric. Note that objects that would capture an observer’s attention are retained while the rest are discarded.

any tracklets ending before the tracklet’s start time, it determines the most similar predecessor tracklet L_k based on the visual similarity score. This is an initial guess for the best predecessor tracklet, however this may belong to the same object, but not be the direct predecessor if there is another tracklet between the two that also belongs to the same object. We check if there are any compatible tracklets between L_i and L_k , which have tracklet L_k as their predecessor. Choosing one of these tracklets as the predecessor would result in L_k still being an earlier predecessor. However, we only wish to choose one of these tracklets if it is the most visually similar tracklet to tracklet L_i (except for L_k). We repeat this procedure iteratively until there are no more tracklets between the L_i its current predecessor L_k , or another tracklet which does not have L_k as its predecessor is the most visually similar tracklet not in the current set of predecessors.

We now have a forest of track hypotheses, each tracklet may only have a single predecessor, but a tracklet can be the predecessor for multiple successor tracklets. The resulting forest has at least one tree whose root node corresponds to one of the tracklets with the earliest starting time. Each path H_i through the trees from a root node to a leaf node in this forest is a possible long-term object track. We now cut this forest into a set of paths which is the best possible set of object tracks by applying a greedy recursive track selection strategy. This selects an optimal path from the forest, which is then added to the set of final long-term object tracks F . This path is then cut from forest, with all nodes belonging to this path being removed from the forest, and the forest rearranging itself into a new set of trees with the remaining nodes. This can be seen in Algorithm 1 part 2 and Figure 4 parts B-E. To score paths we use a combination of a Visual Consistency Score C_i^V and a Temporal Density Score C_i^T .

The visual consistency score is the minimum visual similarity embedding distance between any two tracklets in a path. The temporal density score is the fraction of frames of a video where there is a segment present. This penalizes large temporal gaps between tracklets, making it more likely that objects undergoing short occlusion are correctly tracked, and ensures that the most salient objects are grouped consistently throughout the video, as objects to be

		Ours	VSD [41]	KIS [7]	RVOS [29]	
U17 T-C	$\mathcal{J}\&\mathcal{F}$	Mean	56.4	56.2	51.6	-
		Mean	53.4	53.5	48.7	-
		Recall	60.9	61.3	55.1	-
	\mathcal{F}	Decay	1.5	-2.1	4.0	-
		Mean	59.4	59.0	54.5	-
		Recall	64.1	63.2	59.4	-
U17 T-D	$\mathcal{J}\&\mathcal{F}$	Decay	5.8	0.1	7.7	-
		Mean	58.0	56.5	54.2	22.5
		Mean	54.0	51.7	50.0	17.7
	\mathcal{F}	Recall	62.9	59.9	58.9	16.2
		Decay	3.5	21.7	8.4	1.6
		Mean	62.0	61.4	58.3	27.3
U17 Val	$\mathcal{J}\&\mathcal{F}$	Recall	66.6	65.7	62.1	24.8
		Decay	6.6	15.7	11.4	1.8
		Mean	67.9	56.6	59.9	41.2
	\mathcal{F}	Mean	66.4	51.7	-	36.8
		Recall	76.4	-	-	40.2
		Decay	-0.2	-	-	0.5
\mathcal{F}	Mean	69.3	61.4	-	45.7	
	Recall	76.9	-	-	46.4	
	Decay	0.01	-	-	1.7	

Table 1. Our results compared to all other UVOS methods on the DAVIS 2017 unsupervised benchmarks: `test-challenge` (U17 T-C), `test-dev` (U17 T-D), and `val` (U17 Val). VSD [41] and KIS [7] obtained second and third place (after our method) in the 2019 DAVIS Unsupervised VOS Challenge.

tracked in UVOS are present in mostly all frames. The final path score is a weighted sum of 90% the temporal density and 10% the appearance consistency.

We select paths with the highest score through the forest, add these to a final list of tracks F , and cut these from the current forest, reshaping the forest into a new set of trees. This algorithm select a set of object tracks which do not include any overlapping tracklets and have long-term temporal consistency.

Final Tracks Selection. In UVOS algorithms are not penalized for making predictions that do not overlap with ground-truth. However, the number of total object tracks that can be predicted is still limited. In the DAVIS 2017 unsupervised benchmark this is limited to 20 predictions over the whole video.

UnOVOST predicts a potentially large number of tracks, therefore a final object video saliency estimation step is performed to estimate the 20 most salient objects over the whole video to report for evaluation.

Our video saliency score $S_{sal,i}$ is calculated for each track using each tracklet t_j in the track i :

$$S_{sal,i} = \sum_j temp(t_j) conf(t_j) \quad (1)$$

where $temp(t_j)$ is the temporal length of tracklet j and $conf(t_j)$ is the average of the confidence scores of proposals in tracklet t_j . This video saliency metric prefers tracks

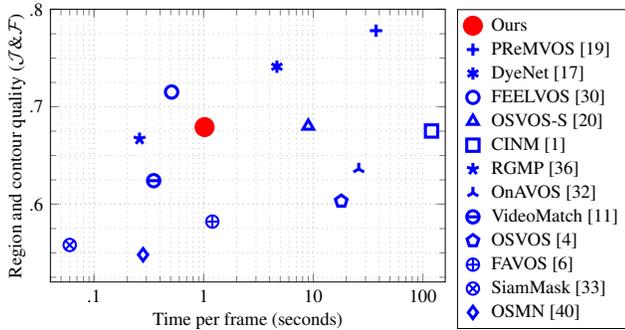


Figure 6. Quality versus timing plot comparing UnOVOST to state-of-the-art semi-supervised methods on DAVIS17 *val*. All methods other than ours are “semi-supervised” and use the given first-frame ground-truth. Our methods obtains similar results while working in an “unsupervised” manner without using any given information about which objects should be tracked.

that are present in many frames and that have high objectness confidence. An example of this is shown in Figure 5.

4. Experiments

Unsupervised VOS Evaluation. We evaluate UnOVOST on the DAVIS 2017 Unsupervised dataset [5]. This contains videos in four sets, with 60 *train*, 30 *val*, 30 *test-dev* and 30 *test-challenge* videos. The *train* and *val* sets contain the same videos as the DAVIS 2017 semi-supervised dataset, however they have been re-annotated according to the definition of the UVOS task. The *test-dev* and *test-challenge* sets contain new videos. All of these datasets include multiple objects per video sequence. Methods are ranked on the $\mathcal{J}\&\mathcal{F}$ metric which is the average of an area overlap (\mathcal{J}) and a boundary overlap (\mathcal{F}) metric. More details can be found in [24].

Table 1 shows our results on these three UVOS benchmarks, and compares our method to three other methods that have presented UVOS results. UnOVOST outperforms all other previous UVOS algorithms over all datasets, often by a large margin. The *val* set is significantly easier than the other two datasets, in this easier setting UnOVOST has the largest margin over the other methods. This shows that when scenes are not too crowded or complex our method does an exceptional job of successfully tracking and segmenting objects through videos. The *test-dev* set is significantly more difficult, and yet the UnOVOST algorithm can still perform extremely well, especially when compared to the performance of RVOS [29]. We present additional qualitative results in Figure 8.

Comparison to Semi-Supervised VOS Methods. As well as comparing to other UVOS methods, we also compare our results on the DAVIS 2017 *val* set to the current state-of-the-art semi-supervised VOS methods. Figure 6 plots the $\mathcal{J}\&\mathcal{F}$ metric of approaches against their runtime per frame. Although these methods use the given first-frame

	U17 Val	U17 T-D	U17 T-C
Mask R-CNN	0.74	0.78	0.77
Optical Flow	0.10	0.14	0.12
Appearance Embedding	0.10	0.15	0.11
UnOVOST Tracking	0.08	0.07	0.06
Total	1.02	1.15	1.06

Table 2. Runtime analysis of UnOVOST on the DAVIS 2017 Unsupervised *val*, *test-dev* and *test-challenge* datasets. Times are seconds per frame.

mask, so they know exactly what objects need to be tracked in the video, our UnOVOST algorithm outperforms many of these methods, even though it operates completely unsupervised without having access to the first frame. Furthermore, many of these methods [4, 19, 32] extensively finetune their segmentation and tracking algorithms on the appearance of the given first-frame objects. Our method still outperforms many of these methods while being significantly faster.

Runtime Analysis. In Table 2 we provide detailed runtime analysis of our UnOVOST algorithm across the three datasets that we test on. The whole algorithm is able to run at around 1 frame per second (fps). The bottleneck of our algorithm is the proposal generation using Mask R-CNN which is more than 70% of the total runtime. We use two other networks to extract features for matching these proposals over time. Our optical-flow and appearance embedding networks, while both reasonably fast at around 10 fps each, make up a combined 20% of the runtime. The remaining runtime for actually running our algorithm, including proposal sub-selection and clipping, tracklet generation, tracklet merging with the FPC algorithm, and track saliency estimation and selection runs in around 0.07 seconds per frame, or at around 15 fps.

Ablation of the Method. We perform an extensive ablation of all design decisions for UnOVOST, for which the results can be found in Figure 7. For all design decisions we use the method which performs best on the training set, even though this is often not the best on the other data splits, except for limiting the output to 20 objects which is required by the evaluation. Interestingly across all experiments, results on *train* and *val* are very similar, whereas *test-dev* is much harder, and often shows different trends in the results than the other two sets.

First we ablate different Mask R-CNN threshold values for our input proposals. With too small or too large a threshold performance degrades significantly. We then ablate the threshold for which to remove proposals if their masks overlap. Again it is important to select a reasonable threshold. Next we ablate the IoU threshold required for merging proposals into tracklets, which is not so important for the easier validation set, but has large effects on the test set results. We use two different tracklet merging strategies, either us-

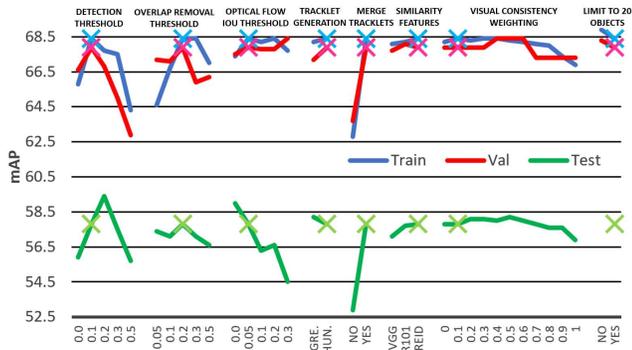


Figure 7. Results of ablating a number of design decisions for UnOVOST on the three splits of the DAVIS 2017 Unsupervised dataset. A cross indicates the selected option chosen as the best performing option on the training set.

ing the Hungarian algorithm for associating proposals into tracklets or alternatively using greedy merging. Hungarian performs better on the train and validation splits but worse on the test split. We show resulting scores when evaluating just our tracklets from the first stage, compared to merging these tracklets in a second stage. Across all sets the second stage is incredibly important. We ablate the use of different similarity features for comparing tracklets, as well as the ReID vectors we compare to using last layer activations of pretrained ResNet 101 [10] and VGG [25] models trained on ImageNet. The ReID vectors outperform the other similarity features on the train and test set, but the ResNet 101 features perform slightly better on the validation set. We ablate different weightings of the Visual Consistency Score and the Temporal Density Score used in our algorithm. Interestingly, the 90:10 ratio which we use as it slightly outperforms other ratios on the training set, is far from the optimal weighting on the other sets. A 50:50 weighting of the two scores performs the best on both the validation and testing set, with more heavily relying on either of scores performing worse. Finally we test the result of our algorithm if we relax the constraint that we can only select 20 objects for evaluation (we could not do this for the hidden test test). Our algorithm performs slightly better in this setting indicating that this restriction removes some correct objects. We recommend that UnOVOST to be used without this restriction step when being applied in the wild, and consider this only an adaption to the dataset and evaluation.

Extension to Video Instance Segmentation. The task of Video Instance Segmentation (VIS) is very similar to UVOS, however in VIS the objects to be tracked must be classified into a set of predefined classes rather than just being salient throughout a video. To investigate the generalization of UnOVOST beyond the UVOS domain we run our algorithm on the YouTube-VIS dataset [39] after training our detector and segmentor on the set of 40 classes in this dataset and adding another classification network to im-



Figure 8. Additional qualitative results of UnOVOST.

prove classification results. Apart from that we run UnOVOST with exactly the same parameters as for the unsupervised DAVIS task. Details of how we trained our detector and segmentor for VIS, and of the classifier we used can be found in the supplemental material.

The VIS task is evaluated using the mAP metric. This is similar to the mAP metric used for instance segmentation [18], however it has been extended from single images to video. Details of mAP for VIS can be found in [39].

The previous state-of-the-art VIS method is MaskTrack R-CNN [39], which achieves a mAP scores of 30.3 and 32.2 on the YouTube-VIS validation and test set respectively. UnOVOST significantly outperforms this, achieving mAP scores of 44.8 and 46.7 on the two benchmarks respectively. With these scores UnOVOST also won the 2019 YouTube-VIS Challenge on Video Instance Segmentation, outperforming 18 other methods. In the supplementary material we present a table comparing results to all previous benchmarks and 2019 challenge entries.

5. Conclusion

In this paper, we present the novel UnOVOST (Unsupervised Offline Video Object Segmentation and Tracking) algorithm for tackling the unsupervised video object segmentation task. Our algorithm is able to track and segment a huge variety of objects in complex scenes by combining both spatio-temporal consistency and visual similarity cues in a novel tracklet based Forest Path Cutting algorithm for performing data association. UnOVOST outperforms all previous UVOS methods, while even performing competitively with many semi-supervised video object segmentation algorithms without requiring any human input as to which objects should be tracked and segmented.

Acknowledgments: This project has been funded, in parts, by ERC Consolidator Grant DeeViSe (ERC-2017-COG-773161) and by a Google Faculty Research Award. We would like to thank Paul Voigtlaender for helpful discussions.

References

- [1] L. Bao, B. Wu, and W. Liu. CNN in MRF: video object segmentation via inference in a cnn-based higher-order spatio-temporal MRF. In *CVPR*, 2018.
- [2] P. Bideau and E. Learned-Miller. A detailed rubric for motion segmentation. *arXiv preprint arXiv:1610.10033*, 2016.
- [3] P. Bideau, A. RoyChowdhury, R. R. Menon, and E. Learned-Miller. The best of both worlds: Combining cnns and geometric constraints for hierarchical motion segmentation. In *CVPR*, 2018.
- [4] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-shot video object segmentation. In *CVPR*, 2017.
- [5] S. Caelles, J. Pont-Tuset, F. Perazzi, A. Montes, K.-K. Maninis, and L. Van Gool. The 2019 davis challenge on vos: Unsupervised multi-object segmentation. *arXiv:1905.00737*, 2019.
- [6] J. Cheng, Y.-H. Tsai, W.-C. Hung, S. Wang, and M.-H. Yang. Fast and accurate online video object segmentation via tracking parts. In *CVPR*, 2018.
- [7] D. Cho, S. Hong, S. Kang, and J. Kim. Key instance selection for unsupervised video object segmentation. *The 2019 DAVIS Challenge on Video Object Segmentation - CVPR Workshops*, 2019.
- [8] A. Dave, P. Tokmakov, and D. Ramanan. Towards segmenting anything that moves. *arXiv 1902.03715*, 2019.
- [9] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017.
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CVPR*, 2016.
- [11] Y.-T. Hu, J.-B. Huang, and A. G. Schwing. Videomatch: Matching based video object segmentation. In *ECCV*, 2018.
- [12] S. D. Jain, B. Xiong, and K. Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. *CVPR*, 2017.
- [13] C. Kim, F. Li, A. Ciptadi, and J. M. Rehg. Multiple hypothesis tracking revisited. *ICCV*, 2015.
- [14] L. Leal-Taixé, A. Milan, K. Schindler, D. Cremers, I. Reid, and S. Roth. Tracking the trackers: an analysis of the state of the art in multiple object tracking. *arXiv preprint arXiv:1704.02781*, 2017.
- [15] S. Li, B. Seybold, A. Vorobyov, A. Fathi, Q. Huang, and C.-C. J. Kuo. Instance embedding transfer to unsupervised video object segmentation. *CVPR*, 2018.
- [16] S. Li, B. Seybold, A. Vorobyov, X. Lei, and C.-C. Jay Kuo. Unsupervised video object segmentation with motion-based bilateral networks. In *ECCV*, September 2018.
- [17] X. Li and C. Change Loy. Video object segmentation with joint re-identification and attention-aware mask propagation. In *ECCV*, 2018.
- [18] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [19] J. Luiten, P. Voigtlaender, and B. Leibe. PRemVOS: Proposal-generation, refinement and merging for video object segmentation. In *ACCV*, 2018.
- [20] K.-K. Maninis, S. Caelles, Y. Chen, J. Pont-Tuset, L. L. Taixé, and L. Van Gool. Video object segmentation without temporal information. *PAMI*, 2018.
- [21] P. Ochs, J. Malik, and T. Brox. Segmentation of moving objects by long term video analysis. *TPAMI*, 2014.
- [22] A. Ošep, W. Mehner, P. Voigtlaender, and B. Leibe. Track, then decide: Category-agnostic vision-based multi-object tracking. *ICRA*, 2018.
- [23] A. Ošep, P. Voigtlaender, J. Luiten, S. Breuers, and B. Leibe. Large-scale object mining for object discovery from unlabeled video. *ICRA*, 2019.
- [24] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. *CVPR*, 2016.
- [25] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [26] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. *arXiv preprint arXiv:1709.02371*, 2017.
- [27] P. Tokmakov, K. Alahari, and C. Schmid. Learning motion patterns in videos. *CVPR*, 2017.
- [28] P. Tokmakov, K. Alahari, and C. Schmid. Learning video object segmentation with visual memory. *ICCV*, 2017.
- [29] C. Ventura, M. Bellver, A. Girbau, A. Salvador, F. Marques, and X. Giro-i Nieto. Rvos: End-to-end recurrent network for video object segmentation. *CVPR*, 2019.
- [30] P. Voigtlaender, Y. Chai, F. Schroff, H. Adam, B. Leibe, and L.-C. Chen. FEELVOS: Fast end-to-end embedding learning for video object segmentation. In *CVPR*, 2019.
- [31] P. Voigtlaender, M. Krause, A. Ošep, J. Luiten, B. B. G. Sekar, A. Geiger, and B. Leibe. MOTs: Multi-object tracking and segmentation. In *CVPR*, 2019.
- [32] P. Voigtlaender and B. Leibe. Online adaptation of convolutional neural networks for video object segmentation. In *BMVC*, 2017.
- [33] Q. Wang, L. Zhang, L. Bertinetto, W. Hu, and P. H. Torr. Fast online object tracking and segmentation: A unifying approach. In *CVPR*, 2019.
- [34] Y. Wu et al. Tensorpack. github.com/tensorpack/, 2016.
- [35] Z. Wu, C. Shen, and A. van den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 2019.
- [36] S. Wug Oh, J.-Y. Lee, K. Sunkavalli, and S. Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, 2018.
- [37] C. Xie, Y. Xiang, Z. Harchaoui, and D. Fox. Object discovery in videos as foreground motion clustering, 2018.
- [38] N. Xu, L. Yang, Y. Fan, J. Yang, D. Yue, Y. Liang, B. Price, S. Cohen, and T. Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *ECCV*, 2018.
- [39] L. Yang, Y. Fan, and N. Xu. Video instance segmentation. *arXiv preprint arXiv:1905.04804*, 2019.
- [40] L. Yang, Y. Wang, X. Xiong, J. Yang, and A. K. Katsaggelos. Efficient video object segmentation via network modulation. In *CVPR*, 2018.

- [41] Z. Yang, Q. Wang, S. Bai, W. Hu, and P. Torr. Video segmentation by detection for the 2019 unsupervised davis challenge. *The 2019 DAVIS Challenge on Video Object Segmentation - CVPR Workshops*, 2019.