

Towards Segmenting Anything That Moves

Achal Dave
Carnegie Mellon University
achald@cmu.edu

Pavel Tokmakov
Carnegie Mellon University
ptokmako@cmu.edu

Deva Ramanan
Carnegie Mellon University
deva@cmu.edu

Abstract

Detecting and segmenting individual objects, regardless of their category, is crucial for many applications such as action detection or robotic interaction. While this problem has been well-studied under the classic formulation of spatio-temporal grouping, state-of-the-art approaches do not make use of learning-based methods. To bridge this gap, we propose a simple learning-based approach for spatio-temporal grouping. Our approach leverages motion cues from optical flow as a bottom-up signal for separating objects from each other. Motion cues are then combined with appearance cues that provide a generic objectness prior for capturing the full extent of objects. We show that our approach outperforms all prior work on the benchmark FBMS dataset. One potential worry with learning-based methods is that they might overfit to the particular type of objects that they have been trained on. To address this concern, we propose two new benchmarks for generic, moving object detection, and show that our model matches top-down methods on common categories, while significantly out-performing both top-down and bottom-up methods on never-before-seen categories.

1. Introduction

People have the remarkable ability to thrive in a staggeringly diverse world, frequently encountering things they have never seen before. Our approaches for machine perception, meanwhile, often remain trapped in a *closed world*, as in the case of object recognition, where approaches are designed to recognize and name one of N pre-defined classes. But practical robot autonomy requires robust perception in the open-world: even a self-driving car must be able to detect never-before-seen obstacles and debris, regardless of what particular semantic *name* it happens to associate with.

In the computer vision community, open-world recognition is typically addressed from a machine-learning perspective such as zero-shot learning [35] or open-set classification [33]. We advocate a different approach that has

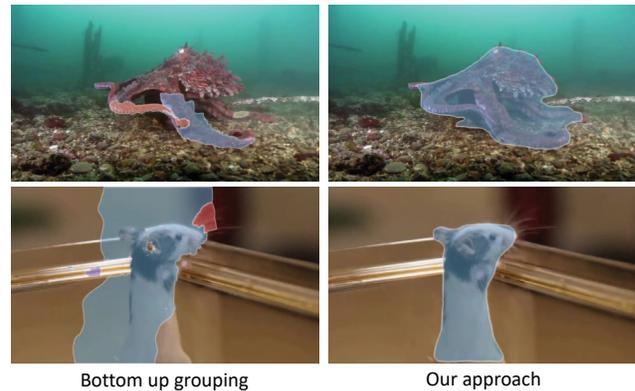


Figure 1: Detecting and segmenting all objects, regardless of category, is key for many perception and robotics tasks. Bottom-up grouping approaches, e.g. [19] (left), aim to tackle this task, but lag behind the quality of closed-world methods that detect a fixed set of N categories. Our work (right) bridges this gap, accurately segmenting generic moving objects, even ones unseen in training.

its roots in classic vision: perceptual grouping. Specifically, we wish to segment out *all* moving object instances in a video stream, including never-before-seen object categories. Defining the notion of a generic, never-before-seen object is notoriously challenging [1]. We intentionally focus on *moving* objects so as to take advantage of the “common fate” principle of grouping: pixels that move together should tend to be grouped together into objects [27].

Indeed, the problem of spatio-temporal grouping is a classic “mid-level” visual understanding task, dating back to the iconic work of Marr [23, 41]. Pre-deep learning solutions tend to follow bottom-up computational strategies for self-organization and clustering, often of long-term pixel trajectories [26, 19]. In the static image case, pixels can be grouped by relying on Gestaltian notions of appearance similarity and curvilinear edge continuity [27]. One long-standing challenge in perceptual organization has been operationalizing these cues into an accurate algorithm for spatio-temporal grouping. *Our key observation is that many of the recent advances in closed-world instance segmentation can be repurposed for open-world spatio-temporal*

grouping.

We first validate the performance of our proposed approach on the Freiburg Berkeley Motion Segmentation benchmark (FBMS). Because the standard measure used in FBMS does not penalize false positives, we find that trivial solutions can score well. We analyze the official metric in detail and propose a new, more informative evaluation. We achieve state-of-the-art results on both measures, and specifically outperform the next-best method of Keuper et al. [19] by 11.4% on our proposed measure.

To further study our method, we introduce the DAVIS-Moving and YTVOS-Moving benchmarks for motion-based grouping. We create these by selecting videos from the DAVIS 2017 [31] and YTVOS [43] datasets where *all* moving objects are labeled. On these new benchmarks, we strongly outperform top-down, closed world methods such as Mask R-CNN, as well as traditional bottom-up grouping methods. In particular, our approach is competitive with a top-down method for categories seen during training, but *outperforms both top-down and bottom-up approaches for unseen categories by 27%*.

To sum up, our contributions are three-fold: (1) we propose the first deep learning-based method for spatio-temporal grouping; (2) we propose a more informative metric and larger, more diverse benchmarks to enable further progress; (3) we report state-of-the-art results on the FBMS dataset and our larger, proposed benchmarks. The code and trained models will be made publicly available.

2. Related Work

Spatio-temporal grouping: Segmenting and tracking objects based on their motion has a rich history. An early work [34] proposed treating this task as a spatio-temporal grouping problem, a philosophy espoused by a number of more recent approaches, including [11, 6, 19], as well as [26], which introduced FBMS. In particular, these methods track each pixel individually with optical flow, encode the motion information of a pixel in a compact descriptor and then obtain an instance segmentation by clustering the pixels based on motion similarity. Unlike these works, our approach is driven primarily by a top-down learning algorithm followed by a simple linking step to generate spatio-temporal segmentations. The most relevant approach in this respect is [9], which trains a CNN to detect (but not segment) moving objects, and combines these detections with clustered pixel trajectories to derive segmentations. By contrast, our approach directly outputs segmentations at each frame, which we link together with an efficient tracker. Very recently, Bideau *et al.* [5] proposed to combine a heuristic-based motion segmentation method [25, 4] with a CNN trained for semantic segmentation for the task of moving object segmentation. Their method, however, does not handle discontinuous motion. In addition, the fact that they rely

strongly on heuristic motion estimates allows our learning-based approach to outperform their method on FBMS by a wide margin. In very recent work, Xie *et al.* [42] introduced a deep learning approach for motion segmentation that segments and tracks moving objects using a recurrent neural network. By comparison, our method uses a simple, overlap-based tracker that performs competitively with the learned tracker from [42] while producing significantly fewer false positive segmentations (see Supplementary).

Foreground/Background Video Segmentation: Several works have focused on the binary version of the video segmentation task, separating all the moving objects from the background. Early approaches [8, 28, 40, 21] relied on heuristics in the optical flow field, such as closed motion boundaries in [28] to identified moving objects. These initial estimates were then refined with appearance, utilizing external cues, such as saliency maps [40], or object shape estimates [21]. Another line of work focused on building probabilistic models of moving objects using optical flow orientations [25, 4]. None of these methods are based on a robust learning framework and struggle to generalize well to unseen videos. The recent introduction of a standard benchmark, DAVIS 2016 [30], has led to a renewed interest. More recent approaches propose deep models for directly estimating motion masks, as in [17, 37, 38]. These approaches are similar to ours in that they also use a two-stream architecture to separately process motion and appearance, but they are unable to segment *individual* object instances, one of our primary goals. Our method separately segments and tracks each individual moving object in a video.

Object Detection: The task of segmenting object instances from still images has seen immense success in recent years, bolstered by large, standard datasets such as COCO [22]. However, this standard task focuses on segmenting every instance of objects belonging to a fixed list of categories, leading to methods that are designed to be blind to objects that fall outside the categories in the training set.

Two recent works have focused on extending these models to detect generic objects. [15] aims to generalize segmentation models to new categories, but requires bounding box annotations for each new category. More relevant to our approach, [18] aims to detect all “object”-like regions in an image, outputting a binary objectness mask. While we share their goal of segmenting unseen objects, our approach additionally provides instance masks for each object.

3. Approach

We propose a two-stream spatio-temporal grouping method that uses appearance and motion cues to segment all moving objects in a video. Our approach, illustrated in Figure 3, takes a frame together with a corresponding optical flow as input, and passes them through an “appearance stream” (top) and a “motion stream” (bottom) respectively.

The resulting features are combined and passed to the joint region proposal network (RPN), which learn to detect and segment moving objects irrespective of their category.

Our approach shares inspiration with prior work that proposes two-stream approaches for object detection [12, 29, 10, 9], with two key differences. First, we design a novel region proposal module that learns to fuse both appearance and motion information to generate moving object detections. Second, to overcome the dearth of appropriate training data, we develop a stage-wise training strategy that allows us to leverage synthetic data to train our motion stream, image datasets to train our appearance stream, and a small amount of real video data to train the joint model.

We first discuss the architecture and training strategy for the motion and appearance streams individually, and then detail how to combine these streams into one coherent architecture. Finally, we describe a simple tracker that we use for linking detections across time, allowing us to produce spatio-temporal groupings that span across many frames.

3.1. Motion-based Segmentation

We start by training a motion-based instance segmentation model. As mentioned above, this requires videos with segmentation masks for all moving objects, which is difficult to obtain. Fortunately, prior work has shown that synthetic data can be used for some low-level tasks, such as flow estimation [7] and binary motion segmentation [37]. Inspired by this, we train our motion stream on the FlyingThings3D dataset [24], which contains nearly 2,700 synthetically generated sequences of 3D objects traveling in randomized trajectories, captured with a camera also traveling along a random trajectory. The dataset provides groundtruth optical flow, as well as segmentations for both static and moving objects (See Figure 2). We train our motion-stream using the moving instance labels from [37], treating all moving objects as a single category, and all other pixels, including static objects, as background. The resulting model learns to segment moving objects irrespective of their category. In fact, this model is oblivious to the whole notion of an object and is capable of segmenting parts that exhibit independent motion (see Figure 5). We discuss more details and variants of this approach in Section 5.3.1.

3.2. Appearance-based Segmentation

In order to incorporate appearance information, we next train an image-based object segmentation model that aims to segment the full extent of generic objects. Fortunately, large datasets exist for training image-based instance segmentation models. Here, we train on the MS COCO dataset [22], which contains approximately 120,000 training images with instance segmentation masks for each object in 80 categories. We could train our appearance stream following the standard Mask R-CNN training procedure,

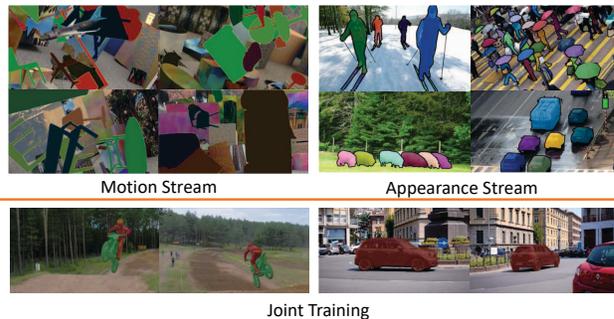


Figure 2: We train our motion stream on FlyingThings3D [24] (top left), our appearance stream on COCO [22] (top right), and our joint model on DAVIS’16 [30] and a YTVOS [43] subset (bottom).

which jointly localizes and classifies each object in an image belonging to the 80 categories. However, this results in a model that, while proficient at segmenting 80 categories, is blind to objects from any other, novel category. Instead, we train an “objectness” Mask R-CNN by combining each of the 80 categories into a single “object” category. In Section 5.3.2, we will show that this “objectness” training (1) provides a significant improvement over standard training, and (2) leads to a model that generalizes surprisingly well to objects that are not labeled in MS COCO.

3.3. Two-Stream Model

Equipped with the individual appearance and motion streams, we now propose a two-stream architecture for fusing these information sources. In order to clearly describe our two-stream model, we take a brief detour to describe the Mask R-CNN architecture. Mask R-CNN contains three stages: (1) **Feature extraction:** a “backbone” network, such as ResNet [14], is used to extract features from an image. (2) **Region proposal:** A region proposal layer uses these features to select regions likely to contain an object. Finally, (3) **Regression:** for each proposed region, the corresponding backbone features are pooled to a fixed size, and fed as input to bounding box and mask regression heads.

To build a two-stream instance segmentation model, we extract the backbone from our individual appearance-based and motion-based segmentation models. Next, as depicted in Figure 3, we propose a “two-stream” RPN that uses these two backbones, instead of a single backbone, to predict proposals from *spatio-temporal* features, extracted from the optical flow (blue) and RGB (orange) backbones. These features are concatenated and fed to a short series of convolutional layers to reduce the dimensionality to match that of Mask R-CNN, allowing us to maintain the architecture of stages (2) and (3). Intuitively, we expect the appearance stream to behave as a generic object detector, and our motion stream to help detect novel objects that the appearance stream may miss and filter out static objects.

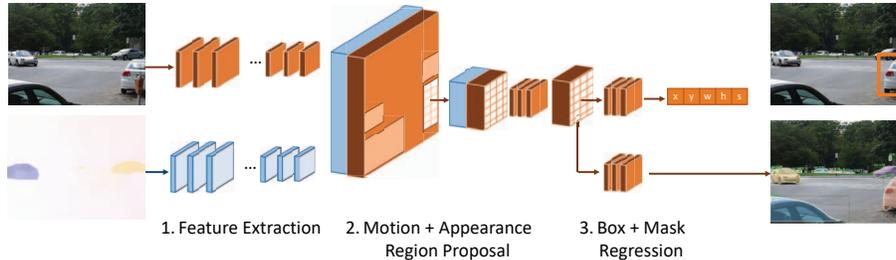


Figure 3: Our model uses an appearance stream (blue) and a motion stream (orange) to extract features from RGB and optical flow frames, respectively. Our region proposal network fuses features from both streams and passes them to the box and mask regression heads.

Although this may appear similar to prior approaches for building a two-stream detection model, it differs in a key detail: prior approaches obtain region proposals either only from appearance features [12, 10, 9], or from appearance and motion features individually [29]. By contrast, we propose a novel proposal module that *learns* to fuse motion and appearance features to find object-like regions.

We train our joint model on subsets of the DAVIS and YouTube Video Object Segmentation datasets (as detailed in Section 5.1). We experiment with various strategies for training this joint model in Section 5.3.3.

3.4. Tracking

So far, we have focused on segmenting moving objects in each frame of a video. To maintain object identities and to continue segmenting objects after they stop moving, we implement a simple, overlap-based tracker inspired by [2]. First, we remove all detections with score below α_{low} . On the first frame, all high scoring detections (score $> \alpha_{\text{high}}$) are used to initialize a track, which we define simply as a sequence of linked detections. At each successive frame, we compute the mask intersection over union between the most recent segmentation for each active track and predicted objects at $t + 1$, and use Hungarian Matching to assign predicted objects to tracks. Unmatched predictions are discarded if their score is $< \alpha_{\text{high}}$; else, they are used to initialize a new track. Tracks that have not been assigned a new object for up to t_{inactive} frames are marked as inactive.

Tracking static objects: To continue tracking moving objects when they stop moving, we need to be able to detect static objects. A naïve way to do this is to run the objectness model trained in Section 3.2 in parallel with our two-stream model at every frame. However, this would be computationally expensive. Fortunately, our appearance stream shares the backbone of the objectness model. Thus, we only need to apply the (inexpensive) stages (2) and (3) of the objectness model on the appearance features extracted by our two-stream network. Using this, we can efficiently output a set of moving and static object predictions for each frame in a video. We merge the two outputs by removing any

predicted static object that overlaps with a predicted moving object. We use the same tracker described above, using only moving objects to initialize tracks.

4. Evaluation

To evaluate methods for spatio-temporal grouping, we desire a metric that rewards segmenting and tracking moving objects, but penalizes the detection of static objects or background. While there has been a rich line of prior work related to our goal, standard metrics surprisingly do not satisfy these criterion. We propose a novel metric that does.

The default metric in FBMS [26] was designed for grouping-based approaches, but does not penalize false positive predictions. Recently, Bideau *et al.* [3] tackled this issue by measuring the difference between the number of groundtruth moving objects and the number of predicted moving objects (ΔObj). However, this complicates method comparisons by relying on two separate metrics; instead, we propose a single and intuitive F-measure that evaluates a method’s ability to detect all and only moving objects.

Figure 4 (middle) visualizes the default FBMS metric which matches each predicted segment with a groundtruth segment so as to maximize IoU overlap, *ignoring* any unmatched predictions. This means the default F-measure does not penalize false positive segments, unfairly favoring methods that generate a large number of predictions. By contrast, our proposed F-measure, depicted in Figure 4 (right), counts unmatched predictions as false positives.

More precisely, we describe our metric roughly following the notation in [26]. For each video, let c_i be the pixels belonging to a predicted region i , and g_j be all the pixels belonging to a groundtruth non-background region j . While [26] omits unlabeled pixels from evaluation, we include all pixels in the groundtruth.

Let P_{ij} be the precision, R_{ij} be the recall, and F_{ij} be the F-measure corresponding to this pair of predicted and groundtruth regions, as follows:

$$P_{ij} = \frac{|c_i \cap g_j|}{|c_i|}, R_{ij} = \frac{|c_i \cap g_j|}{|g_j|}, F_{ij} = \frac{2P_{ij}R_{ij}}{P_{ij} + R_{ij}}$$

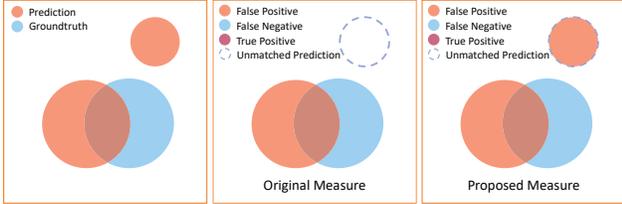


Figure 4: Left: we visualize a toy example with two predicted (red) segmentations and one groundtruth (blue) segmentation. While the original FBMS measure (middle) ignores predicted segments that do not match a groundtruth segment, such as the dashed circle, our proposed measure (right) penalizes all false-positives

Following [26], we use the Hungarian algorithm to find a matching between predictions and groundtruth that maximizes the sum of the F-measure over all assignments. Let $g(c_i)$ be the groundtruth matched to each predicted region; for any c_i that is not matched to a groundtruth cluster, $g(c_i)$ is set to an empty region. We define our metric as follows:

$$P = \frac{\sum_i |c_i \cap g(c_i)|}{\sum_j |c_i|}, R = \frac{\sum_i |c_i \cap g(c_i)|}{\sum_i |g_j|}, F = \frac{2PR}{P + R}$$

Any unlabeled pixel in a predicted region c_i will reduce precision and F-measure, penalizing the segmentation of static or unlabeled objects. In our experiments, we report results with both the official and our proposed measure.

5. Experiments

We first analyze each component of our proposed model with experimental results. Next, we compare our approach to prior work in spatio-temporal grouping on three datasets.

5.1. Datasets

An ideal dataset for training our model would contain a large number of videos where every moving object has labeled instance masks, and static objects are not labeled. Three candidate datasets exist for this task: YouTube Video Object Segmentation (YTVOS) [43], DAVIS 2016 [30], and FBMS [26]. While YTVOS contains over 3,000 short videos with instance segmentation labels, not all objects in these videos are necessarily labeled, and both moving as well as static objects may be labeled. The DAVIS 2016 dataset contains instance segmentation masks (provided with DAVIS 2017) for only the moving objects, but only contains 30 training videos. Finally, although FBMS contains a total of 59 sequences with labeled instance segmentation masks for moving objects, prior work evaluates on the entire dataset, preventing us from training on any sequences in the dataset in order to provide a fair comparison.

To overcome this lack of data, we use heterogeneous data sources to train our model in a stagewise fashion. As described earlier, we train our appearance stream on

COCO [22]. We train our motion stream on FlyingThings3D [24], a synthetic dataset of 2,700 videos of randomly moving 3D objects. Finally, we fine-tune our joint model on DAVIS2016 and the training subset of YTVOS-Moving. We use a held-out set of 100 YTVOS-Moving sequences for evaluation.

5.2. Implementation Details

Network Architecture: Our two-stream model is built off Mask R-CNN [13] with a ResNet-50 backbone. We will publicly release the code and exact configuration for training, highlight some important details here, and note further details in supplementary. All our models are trained using the publicly available PyTorch implementation of Detectron [39]. In general, we use the original hyper-parameters provided by the authors of Mask R-CNN. The backbone for every model is pre-trained on ImageNet [32]. When constructing our two-stream model, we initialize the bounding box and mask heads from the appearance-only model.

Tracking: We set the confidence threshold for initializing tracks, as described in Section 3.4, to $\alpha_{\text{high}} = 0.9$, and remove any detections with confidence lower than $\alpha_{\text{low}} = 0.7$. We allow tracks to stay alive for up to $t_{\text{inactive}} = 10$ frames (approximately 0.33s for most videos), although we found the final results are fairly insensitive to this parameter. To detect objects before they move, we first run our tracker forwards, and then backwards in time.

5.3. Ablation analysis

Evaluation: We analyze our model by benchmarking various configurations on the DAVIS 2016 dataset [30]. For ablation, we found it helpful to use the standard detection mean average precision (mAP) metric [22] in place of video object segmentation metrics, which require tracking and obfuscate analysis of our architecture choices. We report both detection and segmentation mAP at an IoU threshold of 0.5

5.3.1 Motion stream

To begin, we explore training strategies for the motion stream of our model. We train our motion stream on the FlyingThings3D dataset, as described in Section 3.1. This dataset provides groundtruth flow, which we could use for training. However, at inference time, we only have access to noisy, estimated flow. In order to match flow in the real world, we estimate flow on FlyingThings3D using two optical flow estimation methods: FlowNet2 and LiteFlowNet. For both methods, we use the version of their model that is trained on synthetic data and fine-tuned on real data.

In Table 1, we compare three strategies for training on FlyingThings3D. We start by training using only FlowNet2 flow as input (“FlowNet2”). We hypothesize that training directly on noisy, estimated flow can lead to difficul-

Flow type	Det @ 0.5	Seg @ 0.5
FlowNet2	40.5	23.9
FlowNet2 \leftarrow Groundtruth	43.2	24.1
LiteFlowNet \leftarrow Groundtruth	33.8	24.0

Table 1: Comparing training with different flow estimation methods on FlyingThings3D, reporting mAP on DAVIS ’16 val. “ \leftarrow Groundtruth” means we first train with groundtruth (synthetic) flow. See Section 5.3.1 for details.



Figure 5: Despite being trained for segmentation only on synthetic data, our motion stream (visualized) is able to separately segment object parts in real objects. See Sec. 5.3.1 for details.

ties in early training. To overcome this, we train a variant starting with groundtruth flow, and fine-tune on FlowNet2 flow (“FlowNet2 \leftarrow Groundtruth” row). We find that this provides a significant improvement (2.7%). We also considered using a more recent flow estimation method, LiteFlowNet [16] (“LiteFlowNet \leftarrow Groundtruth” row). Surprisingly, we find that FlowNet2 provides significant improvements for detection, despite performing worse on standard flow estimation benchmarks. Qualitatively, we found that FlowNet2 provides sharper results along boundaries than LiteFlowNet, which may aid in localizing objects.

Figure 5 shows qualitative results of the motion stream. Despite never having seen real images with segmentation labels, this model is able to group together parts that move alike, while separating objects with disparate motion.

5.3.2 Appearance Stream

While our motion stream is proficient at grouping similarly-moving pixels, it lacks any priors for real world objects and will not hesitate to oversegment common objects, such as the man in Figure 5. To introduce these useful priors, we turn our attention to the appearance stream of our model.

As described in Section 3.2, we train our appearance stream on the COCO dataset [22]. We evaluate two variants of training. First, we train a standard, “class-specific” Mask R-CNN, that outputs a set of boxes and masks for each of the 80 categories in the COCO Dataset. At inference time, we combine the boxes and masks predicted for each category into a single “object” category. Second, we train an “objectness” Mask R-CNN, by collapsing all the categories in COCO to a single category *before* training.

We show results from these two variants in Table 2. Our

COCO Training	Det @ 0.5	Seg @ 0.5
Class-specific	42.0	40.2
Objectness	49.8	48.3

Table 2: Comparison of training our appearance stream with and without category labels on MS COCO (Class-specific and Objectness, respectively), reporting mAP on DAVIS ’16 val. Training without category labels allows the model to generalize beyond the training categories. See also Figure 6 and Section 5.3.2

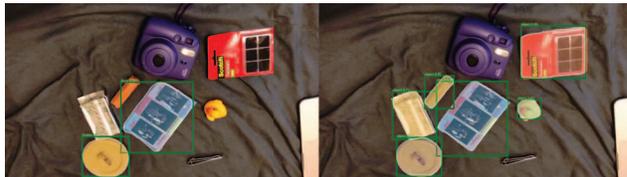


Figure 6: Unlike standard object detectors trained on COCO (left), our objectness model (right) detects objects from categories outside of COCO, such as the packet of film, a roll of quarters, a rubber duck, and a packet of fasteners. Both models visualized at confidence threshold of 0.7. See Section 5.3.2 for details.

“objectness” model significantly outperforms the standard “class-specific” model by nearly 8%. We further compare the two models qualitatively in Figure 6, noting that our objectness model better generalizes to non-COCO categories.

5.3.3 Joint training

Finally, we combine our appearance and flow streams in a single two-stream model, depicted in Figure 3 and described in detail in Section 3.3. We experiment with different strategies for training this joint model. Throughout these experiments, we initialize the flow stream with the “FlowNet2 \leftarrow Groundtruth” model from Section 5.3.1, and use the objectness model from Section 5.3.2 to initialize the appearance stream, the box and mask prediction heads, and the RPN. We show the results in Table 3.

We start by training this joint model directly on the DAVIS 2016 training set, which achieves 79.1% mAP. We note that even with joint-training, using the objectness model for initialization provides a significant boost over using a category-specific detector (73.8%). Next, to maintain the generalizability of the objectness model, we also train a variant where we freeze the weights of the appearance stream. This provides nearly a 3% improvement in accuracy. Similarly, to maintain the generic “grouping” nature of the synthetically-trained flow stream, we freeze the flow stream, providing us with an additional 2% improvement.

Finally, we hypothesize that while features from the flow stream are helpful for localizing generic moving objects, appearance information is sufficient for segmentation. We verify this hypothesis by training one last variant where the

Variant	Det @ 0.5	Seg @ 0.5
Joint Training, class-specific	73.8	70.3
Joint Training, objectness	79.1	73.3
+ Freeze appearance	81.9	76.7
+ Freeze motion	83.7	76.4
+ Freeze mask	83.9	77.4

Table 3: Comparing two-stream training strategies, reporting mAP on DAVIS '16 val. Preserving knowledge from the individual streams is critical for good accuracy. See Section 5.3.3 for details.

Joint Training Data	Det @ 0.5	Seg @ 0.5
DAVIS	83.9	77.4
YTVOS-moving	79.9	75.8
DAVIS \leftarrow YTVOS-moving	85.1	77.9

Table 4: Comparing training sources, reporting mAP on DAVIS '16 val. The lack of static objects in 'YTVOS-moving' leads to worse performance, but fine-tuning on DAVIS provides the best model. See Section 5.3.3 for details.

mask head uses only appearance stream features, and freeze its weights to those of the objectness model. Indeed, this provides a modest improvement of 1% in segmentation AP.

Training Data: Next, we train our joint model on YTVOS-Moving (Section 5.1) and show results in Table 4. Unfortunately, this dataset contains very few static objects, causing the model to detect both static and moving objects, leading to a significant (5%) drop in performance. However, fine-tuning this model on the DAVIS 16 training set leads to our best model (DAVIS \leftarrow YTVOS-moving).

5.4. Comparison to prior work

Official FBMS: We first evaluate our method against prior work on the standard FBMS benchmark in Table 5. As discussed in Section 5.1, this metric does not penalize false positive detections. As expected, our appearance stream alone, despite segmenting both static and moving objects, performs best on this metric ('Ours-A'), outperforming all prior work by 6.4% in F-measure on the TestSet, and 2.2% on the TrainingSet¹. For completion, we also report the performance of our joint model ('Ours-J'), which compares favorably to state-of-the-art despite the flawed metric. Our improvements on this metric are likely driven by improvements in segmentation boundaries (see Figure 7).

Proposed FBMS: Finally, we report results on our proposed metric in Table 6. Recall that our proposed metric generally follows the official metric, but additionally penalizes detection of static objects. We compare to all methods from Table 5 whose final results on FBMS were accessible or provided by the authors through personal communication. On this proposed metric, we first note that, as ex-

¹Note that despite the name, we do not use either set for training.

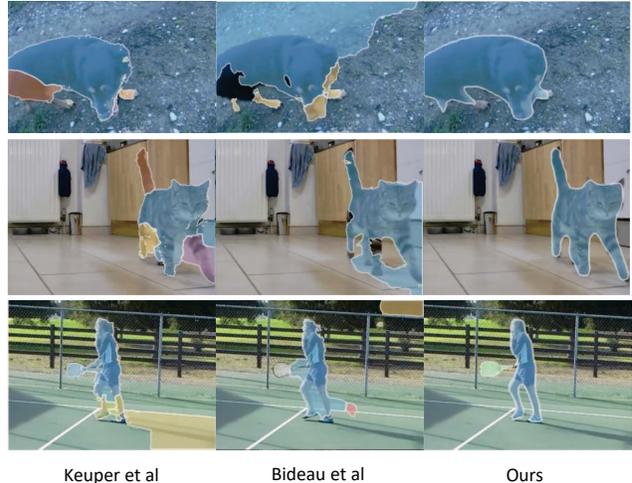


Figure 7: Qualitative results comparing our approach to two state-of-the-art methods. Prior work frequently exhibits over- or under-segmentation, such as the cat (middle row, [19]) and the dog (top row, [5]), respectively. Our method fuses motion and appearance information to segment the full extent of moving objects.

	Training set				Test set			
	P	R	F	N/65	P	R	F	N/69
[36]	83.0	70.1	76.0	23	77.9	59.1	67.3	15
[19]	86.9	71.3	78.4	25	87.6	70.2	77.9	25
[44]	89.5	70.7	79.0	26	91.5	64.8	75.8	27
[20]	93.0	72.7	81.6	29	95.9	65.5	77.9	28
Ours-A	89.2	79.0	83.8	43	88.6	80.4	84.3	40
Ours-J	85.1	78.5	81.7	39	80.8	75.8	78.2	39

Table 5: FBMS 59 results using the official metric [26], which does not penalize detecting unlabeled objects. We report precision (P), recall (R), F-measure (F), and the number of objects for which the F-measure > 0.75 (N). Ours-A is our model's appearance stream only, and Ours-J is our joint model. Both Ours-A and Ours-J out-perform all prior work. As expected, since this metric does not penalize false positives, Ours-A outperforms Ours-J.

pected, the performance of our appearance model baseline is significantly worse than our final, joint model, by 9.2% on TestSet and 6% on TrainingSet in F-measure. More importantly, our final model strongly outperforms prior work in F-measure by 11.3% on the TestSet, and 6.1% on the TrainingSet. In addition to improving segmentation boundaries, our approach effectively removes spurious segmentations of background regions and object parts (Figure 7).

Qualitative results: We qualitatively compare our approach with Keuper *et al.* [19] and Bideau *et al.* [5] in Figure 7². In the top row of Figure 7, [19] oversegments the dog into multiple parts, and [5] merges the dog with the background, whereas our approach fully segments the dog.

²[5] only segments objects while they move. We provide an evaluation using an alternative FBMS labeling they propose in our supplementary.

	Training set			Test set		
	P	R	F	P	R	F
[36]	74.8	61.7	65.5	66.8	49.2	53.6
[19]	68.1	68.5	67.1	70.0	64.6	65.0
Ours-A	61.6	80.4	64.0	66.8	84.7	70.3
Ours-J	75.0	77.8	73.2	77.0	83.0	76.3

Table 6: FBMS 59 results on our proposed metric. Ours-A is our appearance stream, Ours-J is our joint model. We compare to prior methods for which we were able to obtain code or results.

	P	R	F
[19]	39.4	53.8	42.3
Mask R-CNN	70.8	75.6	71.6
Ours	78.3	78.8	78.1

Table 7: DAVIS-Moving results on our proposed metric. We compare to the best FBMS method for which we could obtain code.

	P	R	F
[19] ³	35.3	28.7	26.6
Mask R-CNN	70.4	49.5	53.6
Ours w/o YTVOS	74.5	66.4	68.3

Table 8: YTVOS-Moving results on our proposed metric. For fairness, we evaluate our method *without* YTVOS training. We compare to the best FBMS method for which we could obtain code.

Similarly, the cat in the middle row is over-segmented by [19] and under-segmented by [5], but well-segmented by our approach. In the final row, both [19] and [5] exhibit segmentation and tracking errors; the region corresponding to the man’s foot (colored yellow for Keuper *et al.* and red for Bideau *et al.*) are mistakenly tracked into a background region thus segmenting part of the background as a moving object. Meanwhile, our object-based tracker fully segments the person and the tennis racket with high precision. We show further qualitative results in supplementary material.

DAVIS-Moving: We further evaluate our method on a subset of the DAVIS 17 dataset. Unlike DAVIS 2016, the 2017 version provides instance-level masks for objects, but contains sequences with labeled static or unlabeled moving objects. For evaluation, we manually select 22 of 30 validation videos without these issues, and refer to this subset as DAVIS-Moving. We compare to [19], the best FBMS method we can obtain code for, with our proposed metric in Table 7. Surprisingly, we find a much larger gap in performance on this dataset; while [19] achieves 42.3% on F-measure with our proposed metric, our approach improves significantly to 77.9%. We believe this gap may be due to faster, more articulated motion and higher resolution videos in DAVIS 17, which severely affect [19] but not our method.

³[19] erred on some sequences, so we report numbers on a subset. By comparison, Ours w/o YTVOS achieves 71.9% F-measure on this subset.

	COCO Objects			Novel objects		
	P	R	F	P	R	F
[19]	28.2	25.4	20.6	41.8	31.6	31.9
Mask R-CNN	77.6	60.9	65.1	61.9	37.1	40.6
Ours w/o YTVOS	74.4	66.8	66.8	74.6	66.2	67.6

Table 9: YTVOS-Moving results on seen (COCO) vs. novel objects using our proposed metric.

YTVOS-Moving: Finally, we evaluate on sequences from YTVOS-Moving (selected from YTVOS, as described in Section 5.1). Unlike FBMS and DAVIS, YTVOS contains diverse objects, such as octopuses and snakes. For fairness, we evaluate a version of our final model that was never trained on YTVOS, and show results in Table 8. We show that Mask R-CNN struggles to detect such objects, while our approach strongly improves performance from 53.6% to 67.7% in F-measure. We further break down these results by splitting the YTVOS-Moving dataset into two subsets: videos which contain COCO-category objects, which our model has seen during training, and videos which contain novel objects not from COCO categories in Table 9. While Mask R-CNN is competitive with our approach on COCO categories (underperforming our model by 1.7% F-measure), it significantly underperforms compared to our approach on novel objects, by 27% F-measure. We show qualitative results in supplementary material.

6. Conclusion

We proposed a simple learning-based approach for spatio-temporal grouping. Our method provides two key insights. First, learning based approaches are able to generalize to never-before-seen objects (Section 5.3.2). Second, synthetic data can be used to train a truly generic grouping method with little priors on real world objects. As a result, our approach achieves state-of-the-art results on the FBMS benchmark dataset. Finally, to enable further research in this direction, we introduced a new metric as well as two new benchmarks (DAVIS-Moving, YTVOS-Moving).

Acknowledgements: We thank Pia Bideau for providing evaluation code, Nadine Chang, Kenneth Marino and Senthil Purushwalkam for reviewing drafts and discussions. Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number D17PC00345. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied of IARPA, DOI/IBC or the U.S. Government.

References

- [1] Bogdan Alexe, Thomas Deselaers, and Vittorio Ferrari. What is an object? In *CVPR*. IEEE, 2010. 1
- [2] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Uprocroft. Simple online and realtime tracking. In *ICIP*, 2016. 4
- [3] Pia Bideau and Erik Learned-Miller. A detailed rubric for motion segmentation. *arXiv preprint arXiv:1610.10033*, 2016. 4
- [4] Pia Bideau and Erik Learned-Miller. It’s moving! a probabilistic model for causal motion segmentation in moving camera videos. In *ECCV*, 2016. 2
- [5] Pia Bideau, Aruni RoyChowdhury, Rakesh R Menon, and Erik Learned-Miller. The best of both worlds: Combining CNNs and geometric constraints for hierarchical motion segmentation. In *CVPR*, 2018. 2, 7, 8
- [6] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. In *ECCV*, 2010. 2
- [7] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, 2015. 3
- [8] Alon Faktor and Michal Irani. Video segmentation by non-local consensus voting. In *BMVC*, 2014. 2
- [9] Katerina Fragkiadaki, Pablo Arbelaez, Panna Felsen, and Jitendra Malik. Learning to segment moving objects in videos. In *CVPR*, 2015. 2, 3, 4
- [10] Georgia Gkioxari and Jitendra Malik. Finding action tubes. In *CVPR*, 2015. 3, 4
- [11] Matthias Grundmann, Vivek Kwatra, Mei Han, and Irfan Essa. Efficient hierarchical graph-based video segmentation. In *CVPR*, 2010. 2
- [12] Chunhui Gu, Chen Sun, Sudheendra Vijayanarasimhan, Caroline Pantofaru, David A Ross, George Toderici, Yeqing Li, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, et al. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2017. 3, 4
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 5
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [15] Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, and Ross Girshick. Learning to segment every thing. In *CVPR*, 2018. 2
- [16] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Lite-FlowNet: A lightweight convolutional neural network for optical flow estimation. In *CVPR*, 2018. 6
- [17] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *CVPR*, 2017. 2
- [18] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Pixel objectness. *arXiv preprint arXiv:1701.05349*, 2017. 2
- [19] Margret Keuper, Bjoern Andres, and Thomas Brox. Motion trajectory segmentation via minimum cost multicut. In *ICCV*, 2015. 1, 2, 7, 8
- [20] Naeemullah Khan, Byung-Woo Hong, Anthony Yezzi, and Ganesh Sundaramoorthi. Coarse-to-fine segmentation with shape-tailored continuum scale spaces. In *CVPR*, 2017. 7
- [21] Yong Jae Lee, Jaechul Kim, and Kristen Grauman. Key-segments for video object segmentation. In *ICCV*. IEEE, 2011. 2
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*. Springer, 2014. 2, 3, 5, 6
- [23] David Marr. Vision: A computational investigation into the human representation and processing of visual information. mit press. *Cambridge, Massachusetts*, 1982. 1
- [24] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 2016. 3, 5
- [25] Manjunath Narayana, Allen Hanson, and Erik Learned-Miller. Coherent motion segmentation in moving camera videos using optical flow orientations. In *ICCV*, 2013. 2
- [26] Peter Ochs, Jitendra Malik, and Thomas Brox. Segmentation of moving objects by long term video analysis. *TPAMI*, 36(6):1187–1200, 2014. 1, 2, 4, 5, 7
- [27] Stephen E Palmer. *Vision science: Photons to phenomenology*. MIT press, 1999. 1
- [28] Anestis Papazoglou and Vittorio Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013. 2
- [29] Xiaojiang Peng and Cordelia Schmid. Multi-region two-stream R-CNN for action detection. In *ECCV*. Springer, 2016. 3, 4
- [30] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 2, 3, 5
- [31] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alexander Sorkine-Hornung, and Luc Van Gool. The 2017 DAVIS challenge on video object segmentation. *arXiv:1704.00675*, 2017. 2
- [32] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 5
- [33] Walter J Scheirer, Lalit P Jain, and Terrance E Boult. Probability models for open set recognition. *IEEE TPAMI*, 36(11):2317–2324, 2014. 1
- [34] Jianbo Shi and Jitendra Malik. Motion segmentation and tracking using normalized cuts. In *ICCV*, 1998. 2
- [35] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, 2013. 1
- [36] Brian Taylor, Vasily Karasev, and Stefano Soatto. Causal video object segmentation from persistence of occlusions. In *CVPR*, 2015. 7, 8

- [37] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning motion patterns in videos. In *CVPR*, 2017. 2, 3
- [38] Pavel Tokmakov, Cordelia Schmid, and Karteek Alahari. Learning to segment moving objects. *IJCV*, Sep 2018. 2
- [39] Roy Tseng. Detectron.pytorch. <https://github.com/roytseng-tw/Detectron.pytorch>. 5
- [40] Wenguan Wang, Jianbing Shen, and Fatih Porikli. Saliency-aware geodesic video object segmentation. In *CVPR*, 2015. 2
- [41] Max Wertheimer. Untersuchungen zur lehre von der gestalt. ii. *Psychologische forschung*, 4(1):301–350, 1923. 1
- [42] Christopher Xie, Yu Xiang, Dieter Fox, and Zaid Harchaoui. Object discovery in videos as foreground motion clustering. 2019. 2
- [43] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. YouTube-VOS: A large-scale video object segmentation benchmark. In *ECCV*, 2018. 2, 3, 5
- [44] Yanchao Yang, Ganesh Sundaramoorthi, and Stefano Soatto. Self-occlusions and disocclusions in causal video object segmentation. In *ICCV*, 2015. 7