

Moving Object Segmentation: All You Need Is SAM (and Flow)

Junyu Xie¹, Charig Yang¹, Weidi Xie^{1,2}, and Andrew Zisserman¹

¹ Visual Geometry Group, University of Oxford

² School of Artificial Intelligence, Shanghai Jiao Tong University

{jyx,charig,weidi,az}@robots.ox.ac.uk

<https://www.robots.ox.ac.uk/~vgg/research/flowsam/>

Abstract. The objective of this paper is motion segmentation – discovering and segmenting the moving objects in a video. This is a much studied area with numerous careful, and sometimes complex, approaches and training schemes including: self-supervised learning, learning from synthetic datasets, object-centric representations, amodal representations, and many more. Our interest in this paper is to determine if the Segment Anything model (SAM) can contribute to this task.

We investigate two models for combining SAM with optical flow that harness the segmentation power of SAM with the ability of flow to discover and group moving objects. In the first model, we adapt SAM to take optical flow, rather than RGB, as an input. In the second, SAM takes RGB as an input, and flow is used as a segmentation prompt. These surprisingly simple methods, without any further modifications, outperform all previous approaches by a considerable margin in both single and multi-object benchmarks. We also extend these frame-level segmentations to sequence-level segmentations that maintain object identity. Again, this simple model achieves outstanding performance across multiple moving object segmentation benchmarks.

Keywords: Moving Object Discovery · Video Object Segmentation

1 Introduction

Recent research in image segmentation has been transformative, with the Segment Anything Model (SAM) [13] emerging as a significant breakthrough. Leveraging large-scale datasets and scalable self-labelling, SAM enables flexible image-level segmentation across many scenarios [5, 24, 36, 39, 45, 58], facilitated by user prompts such as boxes, texts, and points. In videos, optical flow has played an important and successful role for *moving object segmentation* – in that it can (i) discover moving objects, (ii) provide crisp boundaries for segmentation, and (iii) group parts of objects together if they move together. It has formed the basis for numerous methods of moving object discovery by self-supervised learning [11, 16, 26, 27, 37, 46, 51]. However, it faces segmentation challenges if objects are momentarily motionless, and in distinguishing foreground objects from

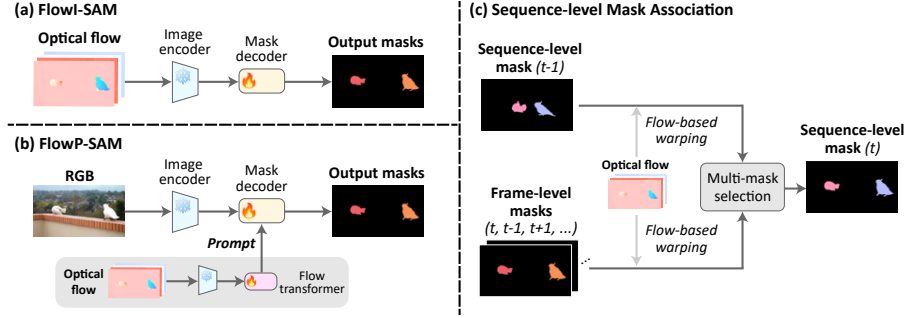


Fig. 1: Adapting SAM for Video Object Segmentation by incorporating flow. (a) **Flow-as-Input:** FlowI-SAM takes in optical flow *only* and predicts frame-level segmentation masks. (b) **Flow-as-Prompt:** FlowP-SAM takes in RGB and applies flow information as a prompt for frame-level segmentation. (c) **Sequence-level mask association:** as a post-processing step, the multi-mask selection module autoregressively transforms *frame-level* mask outputs from FlowI-SAM and/or FlowP-SAM and produces *sequence-level* masks in which object identities are consistent throughout the sequence.

background ‘noise’. This naturally raises the question: “How can we leverage the power of SAM with flow for moving object segmentation in videos?”.

To this end, we explore two simple variants to effectively tailor SAM for motion segmentation. *First*, we introduce FlowI-SAM (Fig. 1a), an adaption of the original SAM that directly processes optical flow as a three-channel *input* image for segmentation, where points on a uniform grid are used as prompts. This approach leverages the ability of SAM to accurately segment moving objects against the static background, by exploiting the distinct textures and clear boundaries present in optical flow fields. However, it has less success in scenes where the optical flow arises from multiple interacting objects as the flow only contains limited information for separating them. *Second*, building on the strong ability of SAM on RGB image segmentation, we propose FlowP-SAM (Fig. 1b) where the input is an RGB frame, and flow guides SAM for moving object segmentation as *prompts*, produced by a trainable prompt generator. This method effectively leverages the ability of SAM on RGB image segmentation, with flow information acting as a selector of moving objects/regions within a frame. Additionally, we extend these methods from frame-level to *sequence-level* video segmentation (Fig. 1c) so that object identities are consistent throughout the sequence. We do this by introducing a matching module that auto-regressively chooses whether to select a new object or propagate the old one based on temporal consistency.

In summary, this paper introduces and explores two models to leverage SAM for moving object segmentation in videos, enabling the principal moving objects to be discriminated from background motions. Our contributions are threefold:

- The FlowI-SAM model, which utilises optical flow as a three-channel input image for precise *per-frame* segmentation and moving object identification.

- The **FlowP-SAM** model, a novel combination of dual-stream (RGB and flow) data, that employs optical flow to generate prompts, guiding SAM to identify and localise the moving objects in RGB images.
- New state-of-the-art unsupervised video object segmentation performance by a large margin on moving object segmentation benchmarks, including DAVIS16, DAVIS17-m, YTVOS18-m, and MoCA.

2 Related Work

Video object segmentation (VOS) is an extensively studied task in computer vision. The objective is to segment the primary object(s) in a video sequence. Numerous benchmarks are developed for evaluating VOS performance, catering to both single-object [17, 19, 30, 33] and multi-object [35, 50] scenarios. Two major protocols are widely explored in VOS research, namely unsupervised [9, 20–23, 41, 51, 54] and semi-supervised VOS [3, 7, 12, 14, 15, 28, 31, 32, 42, 43, 55]. Notably, the term “unsupervised” exclusively indicates that no groundtruth annotation is used *during inference time* (*i.e.*, no inference-time supervision). In contrast, the semi-supervised VOS employs first-frame groundtruth annotations to initiate the object tracking and mask propagation in subsequent frames. This paper focuses on unsupervised VOS and utilises motion as a crucial cue for object discovery.

Motion Segmentation focuses on discovering objects through their movement and generating corresponding segmentation masks. Existing benchmarks for motion segmentation largely overlap with those used for VOS evaluation, especially in the single-object case. For multi-object motion segmentation, datasets [46, 47] have been specifically curated from VOS benchmarks to exclusively focus on sequences with dominant locomotion. There are two major setups in the motion segmentation literature: one that relies on motion information *only* to distinguish moving elements from the background through spatial clustering [26, 27, 51] or explicit supervision [16, 46]; the other [2, 11, 25, 34, 47, 53] that enhances motion-based object discovery by incorporating appearance information. We term these two approaches “flow-only” and “RGB-based” segmentation, respectively, and explore both setups in this work.

Segment Anything Model (SAM) [13] has demonstrated impressive ability on image segmentation across various scenarios. It was trained on the SA-1B datasets with over one billion self-labelled masks and 11 million images. Such large-scale training renders it a strong zero-shot generalisability to unseen domains. Many works adapt the SAM model to perform different tasks, such as tracking [8], change detection [61], and 3D segmentation [4]. Some other works extend SAM towards more efficient models [49, 56, 60], and more domains [5, 24, 36, 45]. However, most studies follow the default prompt options in SAM (*i.e.*, points, boxes, and masks). Recent works [38, 63] have shown that more versatile prompts, including scribbles and visual references, can lead to improvements. In this paper, we explore a novel route that prompts SAM with optical flow and demonstrate its effectiveness for moving object segmentation.




Fig. 2: Overview of FlowI-SAM. (a) Inference pipeline of FlowI-SAM. **(b)** Architecture of FlowI-SAM with trainable parameters labelled. The point prompt token is generated by a frozen prompt encoder.

3 SAM Preliminaries

The Segment Anything Model (SAM) is engineered for high-precision image segmentation, accommodating both user-specified prompts and a fully autonomous operation mode. When guided by user input, SAM accepts various forms of prompts including points, boxes, masks, or textual descriptions to accurately delineate the segmentation targets. Alternatively, in its automatic mode, SAM uses points on a uniform grid as prompts, to propose all plausible segmentation masks that capture objects and their hierarchical subdivisions—objects, parts, and subparts. In this case, the inference is repeated for each prompt of the grid, generating masks for each prompt in turn, and the final mask selection is guided by the predicted mask IoU scores.

Architecturally, SAM is underpinned by three foundational components: (i) Image encoder extracts strong image features via a heavy Vision Transformer (ViT) backbone, which is pre-trained by the Masked Auto-Encoder (MAE) approach; (ii) The prompt encoder converts the input prompts into positional information which helps with locating the segmentation target; (iii) Mask decoder features a light-weight two-way transformer that takes in a combination of encoded prompt tokens, learnable mask tokens, and an IoU prediction token as input queries. These queries iteratively interact with the dense spatial features from image encoder, leading to the final mask predictions and IoU estimations. In the next sections, we describe two distinct, yet simple variants to effectively tailor SAM for motion segmentation.

4 Frame-Level Segmentation I: Flow as Input

In this section, we focus on discovering moving objects from individual frames by exploiting motion information *only*, to yield corresponding segmentation masks. Formally, given the optical flow input $F_t \in \mathbb{R}^{H \times W \times 3}$ at frame t , we aim to predict a segmentation mask $M_t^i \in \mathbb{R}^{H \times W}$ together with a foreground object IoU (fIoU) score $s_{\text{fIoU},t}^i \in \mathbb{R}^1$ for each object i ,

$$\{M_t^i, s_{\text{fIoU},t}^i\}_{i=0}^N = \Phi_{\text{FlowI-SAM}}(F_t) \quad (1)$$

To adapt SAM for this new task, we formulate **FlowI-SAM** ($\Phi_{\text{FlowI-SAM}}$) by finetuning it on optical **Flow Inputs**, and re-purpose the original IoU prediction head to instead predict the floU, as illustrated in Fig. 2b. By definition, floU is a scalar that measures the “objectness”: floU is 0 if the mask belongs to the background, and equal to the IoU between predicted and GT object masks for foreground moving object masks. A high floU indicates the predicted mask corresponds to the entire object, while a low floU might suggest the mask is erroneous or only captures a small part of the object.

Flow Inputs with Multiple Frame Gaps. To mitigate the effect of noisy optical flow, *i.e.* complicated flow fields due to stationary parts, articulated motion, and object interactions, *etc.*, we consider multiple flow inputs $\{F_{t,g}\}$ with different frame gaps (*e.g.*, $g \in \{(1,-1), (2,-2)\}$) for both training and evaluation stages. These multi-gap flow inputs are **independently** processed by the image encoder to obtain dense spatial features $\{d_{t,g}\}$ at a lower resolution $h \times w$, which are then combined by averaging the spatial feature maps across different flow gaps, *i.e.*, $d_t = \text{Average}_g(\{d_{t,g}\}) \in \mathbb{R}^{h \times w \times d}$. These averaged spatial features are then treated as keys and values in the mask decoder.

FlowI-SAM Inference. To discover all moving objects from flow input, the FlowI-SAM model is prompted by points on a uniform grid. Each point prompt outputs a pair of mask and objectness score predictions. This mechanism is the same as in the original SAM formulation, and is illustrated in Fig. 2a. The final segmentation is selected using Non-Maximum Suppression (NMS) based on the predicted floU and overlap ratio.

FlowI-SAM Training. To adapt the pre-trained SAM model for optical flow inputs, we finetune the lightweight mask decoder, while the image encoder and the prompt encoder remain frozen. The overall loss is formulated as:

$$\mathcal{L}_{\text{FlowI-SAM}} = \frac{1}{NT} \sum_{i,t}^{N,T} \left(\mathcal{L}_{\text{BCE}}(M_t^i, \hat{M}_t^i) + \lambda_f \|s_{\text{floU},t}^i - \hat{s}_{\text{floU},t}^i\|^2 \right) \quad (2)$$

where \hat{M}_t^i and $\hat{s}_{\text{floU},t}^i$ denote the groundtruth segmentation masks and floU, and λ_f is a scale factor.

5 Frame-Level Segmentation II: Flow as Prompt

In this section, we adapt SAM for video object segmentation by processing RGB frames, with optical flow as a prompt. We term this frame-level segmentation architecture **FlowP-SAM** for **Flow** as **Prompt** SAM. As shown in Fig. 3b, FlowP-SAM encompasses two major modules, namely the flow prompt generator and the segmentation module. The flow prompt generator takes optical flow as inputs, and produces flow prompts that can be used as supplemental queries to infer frame-level segmentation masks M_t^i from RGB inputs I_t . Formally,

$$\{M_t^i, s_{\text{floU},t}^i, s_{\text{MOS},t}^i\}_{i=0}^N = \Phi_{\text{FlowP-SAM}}(F_t, I_t) \quad (3)$$

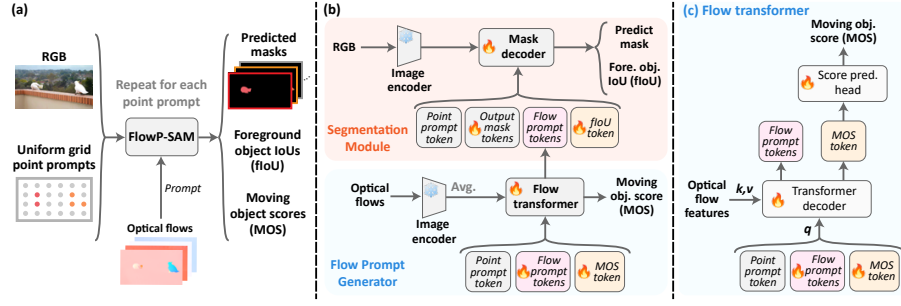


Fig. 3: Overview of FlowP-SAM. (a) Inference pipeline of FlowP-SAM. (b) Architecture of FlowP-SAM. The flow prompt generator produces flow prompts to be injected into a SAM-like RGB-based segmentation module. Both modules take in the same point prompt token, which is obtained from a frozen prompt encoder. (c) Detailed architecture of the flow transformer. The input tokens function as queries within a lightweight transformer decoder, iteratively attending to dense flow features. The output moving object score (MOS) token is then processed by an MLP-based head to predict a score indicating whether the input point prompt corresponds to a moving object.

where $s_{\text{MOS},t}^i$ indicates the moving object score (MOS) predicted by the flow prompt generator, while $s_{\text{floU},t}^i$ denotes the foreground object IoU (floU) estimated by the segmentation module. Specifically, MOS measures whether each input point prompt (therefore the resultant mask) is within a moving object region based on observing flow fields. Groundtruth MOS scores are binary (*i.e.*, $\hat{s}_{\text{MOS},t}^i = 1$ if the point prompt is within GT annotation, and $\hat{s}_{\text{MOS},t}^i = 0$ otherwise). On the other hand, floU follows the same formulation as in FlowI-SAM, *i.e.*, predicting IoUs for foreground objects and yielding 0 for background regions.

Flow Prompt Generator consists of (i) a frozen SAM image encoder, where the dense spatial features are extracted from optical flow inputs at different frame gaps, followed by an averaging across frame gaps; (ii) a flow transformer, with the detailed architecture depicted in Fig. 3c, where we first stack the input point prompt (*i.e.*, a positional embedding) with learnable flow prompts and moving object score (MOS) tokens to form queries. These queries then iteratively attend to dense flow features in a lightweight transformer decoder. There are two outputs from the flow prompt generator, namely, the refined flow prompts, and an MOS token, which is subsequently processed by an MLP-based head to yield a final moving object score.

Segmentation Module has a structure that resembles the original SAM, except for two adaptations: (i) The IoU-prediction head is re-purposed to predict foreground object scores (floU) (same as the FlowI-SAM); (ii) The outputs tokens from flow prompt generator are injected as additional query inputs.

FlowP-SAM Inference. Similar to FlowI-SAM, we prompt FlowP-SAM with single point prompts from a uniform grid to iteratively predict possible segmentation masks, together with MOS and floU estimations. These predicted scores are

averaged, *i.e.*, $(s_{\text{MOS},t}^i + s_{\text{IoU},t}^i)/2$, and then utilised to guide post-processing which involves the NMS and overlay of selected masks.

FlowP-SAM Training. We train FlowP-SAM in an end-to-end fashion while keeping the SAM pre-trained prompt encoder and image encoders frozen. The flow transformer is trained from scratch,

$$\begin{aligned} \mathcal{L}_{\text{FlowP-SAM}} = & \frac{1}{NT} \sum_{i,t}^{N,T} \left(\mathcal{L}_{\text{BCE}}(M_t^i, \hat{M}_t^i) \right. \\ & \left. + \lambda_m \mathcal{L}_{\text{BCE}}(s_{\text{MOS},t}^i, \hat{s}_{\text{MOS},t}^i) + \lambda_f \|s_{\text{IoU},t}^i - \hat{s}_{\text{IoU},t}^i\|^2 \right) \end{aligned} \quad (4)$$

where \hat{M}_t^i corresponds to the groundtruth mask. $\hat{s}_{\text{MOS},t}^i$ and $\hat{s}_{\text{IoU},t}^i$ indicate the groundtruth of two predicted scores, with λ_m and λ_f being the scale factors.

6 Sequence-level Mask Association

In this section, we outline our method for linking the frame-wise predictions for each moving object into a continuous track throughout the sequence. Specifically, we compute two types of masks: *frame-wise* masks M at the current frame using the model (FlowI-SAM and/or FlowP-SAM); and *sequence-level* masks \mathcal{M} , that are obtained by propagating the initial frame prediction with optical flow, we then update the mask of the current frame by making a comparison between the two. The following section details our update mechanism.

Update Mechanism. This process aims to associate object masks across frames, as well as to determine whether the sequence-level results at a particular frame should be obtained directly from frame-level predictions at that frame or by propagating from previous frame results.

Specifically, given a sequence-level mask for object i at frame $t - 1$ (*i.e.*, \mathcal{M}_{t-1}^i), we first warp it to frame t using optical flow,

$$\mathcal{M}_{t \leftarrow t-1}^i = \text{warp}(\mathcal{M}_{t-1}^i, F_{t-1}) \quad (5)$$

We then consider three sets of masks: (i) the warped masks $\{\mathcal{M}_{t \leftarrow t-1}^i\}$; (ii) the frame-level predictions $\{M_t^i\}$ at frame t ; and (iii) the frame-level predictions from neighboring frames (with Δt gap) after aligning them to the current frame using optical flow (*i.e.*, $\{M_{t \leftarrow t+\Delta t}^i\}$). For each pair of mask sets, we perform a pairwise Hungarian matching based on the IoU score, resulting in three pairings in total. The Hungarian matched pairs can then reflect the *temporal consistency* across these predictions based on the *transitivity* principle, *i.e.* if object i in (i) matches with object j in (ii) and object k in (iii), then the latter two objects must also match with each other. If such transitivity holds, we set the consistency score $c_i = 1$, and $c_i = 0$ otherwise.

This matching process is repeated for $\Delta t \in \{1, 2, -1, -2\}$, resulting in an averaged consistency score \bar{c}_i , which guides the following mask update:

$$\mathcal{M}_t^i = \begin{cases} M_t^i & \text{if } \bar{c}_i \geq 0.5 \\ \mathcal{M}_{t \leftarrow t-1}^i & \text{if } \bar{c}_i < 0.5 \end{cases} \quad (6)$$

where \mathcal{M}_t^i denotes the sequence-level mask prediction for object i at frame t .

The rationale behind this is that the two methods have their own strengths and drawbacks: propagation is safe in preserving the object identity, but the mask quality degrades over time, whereas updating ensures high-quality masks but comes with the risk of mis-associating object identities. Thus, if the current frame-wise mask is temporally consistent, then we can be reasonably confident to update the mask, if not, then we choose the safer option and propagate the previous mask.

Note, we do this separately for each object $i \in N$, which gets updated or propagated independently. We lastly layer all objects back together (to its original order) and remove any overlaps to obtain the final sequence-level predictions.

7 Experiments

7.1 Datasets

Single-Object Benchmarks. For single-object motion segmentation, we adopt standard datasets, including DAVIS2016 [33], SegTrackv2 [19], FBMS-59 [29], and MoCA [17]. Although SegTrackv2 and FBMS-59 include a few multi-object sequences, following the common practice [16, 51], we treat them as single-object benchmarks by grouping all moving objects into a single foreground mask. MoCA stands for Moving Camouflaged Animals, designed as a camouflaged object detection benchmark. Following [16, 46, 51], we adopt a filtered MoCA dataset by excluding videos with predominantly no locomotion.

Multi-Object Benchmarks. In terms of multi-object segmentation, we report the performance on DAVIS2017 [35], DAVIS2017-motion [35, 46], and YouTubeVOS2018-motion [47, 50], where DAVIS2017 is characterised by predominantly moving objects, each annotated as distinct entities. For example, a man riding a horse would be separately labelled. In contrast, DAVIS2017-motion re-annotates the objects based on their joint movements such that objects with shared motion are annotated as a single entity. For example, a man riding a horse is annotated as a single entity due to their shared motion.

The YouTubeVOS2018-motion [47] dataset is a curated subset of the original YouTubeVOS2018 [50]. It specifically excludes video sequences involving common motion, severe partial motion, and stationary objects, making it ideally suited for motion segmentation evaluation. Whereas, the original dataset also annotates many stationary objects and only provides partial annotations for a subset of moving objects.

Summary of Evaluation Datasets. To investigate the role of motion in object discovery and segmentation, we adopt all aforementioned benchmarks, which

consist of only predominantly moving object sequences. Notably, for the evaluation of **FlowI-SAM**, we exclude the class-labelled DAVIS17 dataset, as commonly moving objects cannot be separated solely based on motion cues.

Training Datasets. To adapt the RGB pre-trained SAM for moving object discovery and motion segmentation, we train both **FlowI-SAM** and **FlowP-SAM** first on the synthetic dataset introduced by [46], as described below, and then on real-world video datasets, including DAVIS16, DAVIS17, and DAVIS17-m.

7.2 Evaluation Metrics

To assess the accuracy of predicted masks, we report intersection-over-union (\mathcal{J}), except for MoCA where only the ground-truth bounding boxes are given, we instead follow the literature [51] and report the detection success rate (SR). Regarding multi-object benchmarks, we additionally report the contour accuracy (\mathcal{F}) in the arxiv version [48].

In this work, we differentiate between the frame-level and sequence-level methods, and adopt two distinct evaluation protocols: (i) Since frame-level methods generate the segmentation *independently* for each frame, we apply per-frame Hungarian matching to match the object masks between predictions with the groundtruth, before the evaluation; (ii) Conversely, sequence-level methods employ an extra step to link object masks across frames. As a result, the Hungarian matching is conducted globally for each sequence, *i.e.*, the object IDs between predicted and groundtruth masks are matched once per sequence. Given the added complexity and potential errors during frame-wise object association, sequence-level predictions are often considered a greater challenge.

7.3 Implementation Details

In this section, we summarise the experimental setting in our frame-level segmentation models. For more information regarding detailed architectures, hyperparameter settings, and sequence-level mask associations, please refer to the arxiv version [48].

Flow Computation. We adopt an off-the-shelf method (RAFT [40]) to estimate optical flow with multiple frame gaps at (1,-1) and (2,-2), except for YTVOS18-m and FBMS-59, where higher frame gaps at (3,-3) and (6,-6) are used to compensate for slow motion. Following common practice [46, 51], we convert optical flow into 3-channel RGB format using a standard color wheel [1].

Model Settings. For both **FlowI-SAM** and **FlowP-SAM**, we follow the default SAM setting and adopt the first output mask token (out of four) for mask predictions. For **FlowI-SAM**, we deploy two versions of the pre-trained SAM image encoder, specifically ViT-B and ViT-H, to extract optical flow features.

Regarding **FlowP-SAM**, for efficiency reasons, we utilise ViT-B to encode optical flows and employ ViT-H as the image encoder for RGB frames. We initialise the flow prompt generator with 4 learnable flow prompt tokens, which are subsequently processed by a light-weight two-layer transformer decoder in the flow transformer module.

Stage	Predicted scores for post-processing	Flow prompt	FT mask decoder	DAVIS17 $\mathcal{J} \uparrow$	DAVIS16 $\mathcal{J} \uparrow$
	IoU	\times	\times	25.2	30.3
+ Train flow prompt generator	IoU	\checkmark	\times	61.9	80.6
	(MOS+IoU)/2	\checkmark	\times	63.7	81.4
+ Finetune segment -ation module	(MOS+IoU)/2	\checkmark	\checkmark	65.5	81.5
	(MOS+fIoU)/2	\checkmark	\checkmark	69.9	86.1

Table 1: Ablation analysis of FlowP-SAM. The study starts from the vanilla SAM checkpoint and progressively introduces new components (labelled in blue). ‘‘MOS’’ is short for the moving object score, and ‘‘fIoU’’ indicates the foreground object IoU. The results are shown for frame-level predictions.

Evaluation Settings. At inference time, for FlowI-SAM with flow-only inputs, we input independent point prompts over a 10×10 uniform grid, while for FlowP-SAM, to take account for more complicated RGB textures, we consider a large grid size of 20×20 .

Mask Selection over Multiple Point Prompts. During post-processing, we utilise the predicted scores (fIoU for FlowI-SAM, and an average of MOS and fIoU for FlowP-SAM) as guidance throughout the mask selection process: (i) Non-Maximum Suppression (NMS) is applied to filter out repeating masks and keep the ones with higher scores; (ii) The remaining masks are then ranked according to their scores and top- n masks are retained ($n = 5$ for FlowI-SAM, and $n = 10$ for FlowP-SAM); (iii) These n masks are overlaid by allocating masks with higher scores at the front.

Training Settings. The training is performed in two stages, which involves synthetic pre-training on the dataset proposed by [46], followed by finetuning on the real DAVIS sequences, as detailed in the arxiv version [48]. YTVOS is not used for fine-tuning as there is only a low proportion of moving object sequences. We train both models in an end-to-end manner using the Adam Optimiser at a learning rate of $3e^{-5}$. The training was conducted on a single NVIDIA A40 GPU, with each mode taking roughly 3 days to reach full convergence.

7.4 Ablation Study

In this section, we present a series of ablation studies on key designs in the per-frame FlowP-SAM model. We refer the reader to the arxiv version [48] for a more detailed ablation analysis on the designs in FlowI-SAM and FlowP-SAM models, as well as in our sequence-level method.

Ablation Studies for FlowP-SAM. As illustrated in Table 1, we start from the vanilla SAM and progressively add our proposed components. Note that, we adopt the same inference pipeline (*i.e.*, the same point prompts and post-processing steps) for all predictions shown. Since the foreground object IoU (fIoU) is not predicted by the vanilla SAM, we instead apply default IoU predictions to guide the mask selection.

We *train the flow prompt generator* to simultaneously predict flow prompt tokens and moving object scores (MOS). The injection of flow prompts into the standard RGB-based SAM architecture results in notable enhancements, verifying the value of motion information for accurately determining object positions and shapes. Additionally, employing MOS as additional post-processing guidance yields further improvements.

Upon *finetuning the segmentation module*, we observe a slight enhancement in performance. Finally, substituting the default IoU predictions with fIoU scores achieves more precise mask selection, as evidenced by the improved results.

Discussion on the effectiveness of MOS and fIoU. As outlined in Sect. 3, the original SAM framework over-segments images into objects, parts, and sub-parts, which cannot be further distinguished using the default SAM IoU estimations. To adapt this setup for *object-only* discovery, we propose two new scores, *i.e.*, MOS and fIoU, as new criteria to filter out non-object masks. These scores effectively assess the “objectness” of masks: MOS determines if the predicted masks represent moving objects, while fIoU evaluates whether the masks depict complete objects and are not background segments. The effectiveness of this adaptation is validated in Table 1, where replacing IoU estimations with MOS + fIoU scores leads to noticeable performance boosts.

7.5 Quantitative Results

Given the distinct evaluation protocols outlined in Sec. 7.2, we report our method separately, with a frame-level analysis for **FlowI-SAM** (introduced in Sec. 4) and **FlowP-SAM** (introduced in Sec. 5), followed by a sequence-level evaluation.

Frame-Level Performance. Table 2 distinguishes between flow-only and RGB-based methods, where the former adopts optical flow as the only input modality, and the latter takes in RGB frames with optional flow inputs. Note that, the performance for some recent self-supervised methods is also reported, owing to the lack of the supervised baselines.

For flow-only segmentation, our **FlowI-SAM** (with both SAM image encoders) outperforms the previous methods by a large margin ($>10\%$). For RGB-based segmentation, our **FlowP-SAM** also achieves state-of-the-art performance, particularly excelling at multi-object benchmarks. By combining these two frame-level predictions (**FlowI-SAM+FlowP-SAM**), we observe further performance boosts. This suggests the complementary roles of the flow and RGB modalities in frame-level segmentation, particularly when there are multiple moving objects involved. In particular, we show that using both models in tandem by layering **FlowI-SAM**’s predictions behind that of **FlowP-SAM** allows the model to fill in on missed predictions (such as motion blur, poor lighting, or small objects).

Sequence-Level Performance. For flow-based segmentation, we apply the mask association technique introduced in Sec. 6 to obtain sequence-level predictions from per-frame **FlowI-SAM** results. To ensure a fair comparison, we additionally finetune the synthetic-trained OCLR [46] model on the real-world dataset (DAVIS) with groundtruth annotations provided, resulting in “OCLR-real” results. As shown by the top part of Table 3, **FlowI-SAM** (seq) demonstrates

Model	Flow	RGB	Multi-object benchmarks			Single-object benchmarks			
			YTVOS18-m $\mathcal{J} \uparrow$	DAVIS17-m $\mathcal{J} \uparrow$	DAVIS17 $\mathcal{J} \uparrow$	DAVIS16 $\mathcal{J} \uparrow$	STv2 $\mathcal{J} \uparrow$	FBMS $\mathcal{J} \uparrow$	MoCA SR \uparrow
Flow-only methods									
COD [17]	✓	✗	—	—	—	65.3	—	—	0.236
[†] MG [51]	✓	✗	37.0	38.4	—	68.3	58.6	53.1	0.484
[†] EM [26]	✓	✗	—	—	—	69.3	55.5	57.8	—
FlowI-SAM (ViT-B)	✓	✗	56.7	63.2	—	79.4	69.0	72.9	0.628
FlowI-SAM (ViT-H)	✓	✗	58.6	65.7	—	79.1	70.1	75.1	0.625
RGB-based methods									
[†] VideoCutLER [44]	✗	✓	59.0	57.4	41.7	—	—	—	—
[†] Safadoust et al. [37]	✗	✓	—	59.3	—	—	—	—	—
MATNet [62]	✓	✓	—	—	56.7	82.4	50.4	76.1	0.544
DystaB [53]	✗	✓	—	—	—	82.8	74.2	75.8	—
AMC-Net [52]	✓	✓	—	—	—	84.5	—	76.5	—
TransportNet [57]	✓	✓	—	—	—	84.5	—	78.7	—
TMO [10]	✓	✓	—	—	—	85.6	—	79.9	—
FlowP-SAM	✓	✓	76.9	78.5	69.9	86.1	83.9	87.9	0.645
FlowP-SAM+FlowI-SAM	✓	✓	77.4	80.0	71.6	86.2	84.2	88.7	0.645

Table 2: Frame-level comparison on video object segmentation benchmarks. “[†]” indicates models that are trained without human annotations. For the results in the last row, we combine frame-level predictions from FlowP-SAM and FlowI-SAM (ViT-H).

Model	Flow	RGB	Multi-object benchmarks			Single-object benchmarks			
			YTVOS18-m $\mathcal{J} \uparrow$	DAVIS17-m $\mathcal{J} \uparrow$	DAVIS17 $\mathcal{J} \uparrow$	DAVIS16 $\mathcal{J} \uparrow$	STv2 $\mathcal{J} \uparrow$	FBMS $\mathcal{J} \uparrow$	MoCA SR \uparrow
Flow-only methods									
[†] SIMO [16]	✓	✗	—	—	—	67.8	62.0	—	0.566
[†] Meunier et al. [27]	✓	✗	—	—	—	73.2	55.0	59.4	—
[†] OCLR [46]	✓	✗	46.5	54.5	—	72.1	67.6	70.0	0.599
OCLR-real [46]	✓	✗	49.5	55.7	—	73.3	65.9	70.5	0.605
FlowI-SAM (seq, ViT-B)	✓	✗	51.9	60.0	—	78.4	66.9	69.0	0.615
FlowI-SAM (seq, ViT-H)	✓	✗	53.8	61.5	—	78.0	67.7	71.5	0.604
RGB-based methods									
UnOVOST [23]	✗	✓	—	—	66.4	—	—	—	—
Propose-Reduce [21]	✗	✓	—	—	67.0	—	—	—	—
OCLR-flow [46] + SAM [13]	✓	✓	57.0	62.0	—	80.6	71.5	79.2	—
PMN [18]	✓	✓	—	—	—	85.6	—	77.8	—
Xie et al. [47] + SAM [13]	✓	✓	71.1	70.9	—	86.6	81.3	85.7	—
DEVA [6]	✗	✓	—	—	70.4	87.6	—	—	—
UVOSAM [59]	✗	✓	—	—	77.5	—	—	—	—
FlowP-SAM+FlowI-SAM (seq)	✓	✓	74.7	74.3	71.0	87.7	80.1	82.8	0.647

Table 3: Sequence-level comparison on video object segmentation benchmarks. “[†]” indicates models that are trained without human annotations. “seq” indicates that our sequence-level predictions with object masks matched across frames. We adopt FlowP-SAM and FlowI-SAM (ViT-H) to obtain the results in the last row.

superior performance against OCLR-real, benefiting from the robust prior knowledge in pre-trained SAM.

For RGB-based segmentation, we obtain our sequence-level predictions by FlowI-SAM+FlowP-SAM. As shown in the lower part of Table 3, our method




Fig. 4: Qualitative comparison of flow-only segmentation methods on DAVIS (left), YTVOS (middle), and MoCA (right) sequences. Our **FlowI-SAM (seq)** successfully identifies moving objects from noisy optical flow background (*e.g.*, the ducks in the fourth column).




Fig. 5: Qualitative comparison of RGB-based segmentation methods on DAVIS (left), YTVOS (middle), and SegTrackv2 (right). While the previous method (the third row) struggles to disentangle multiple moving objects (*e.g.*, mixed gold fishes in the second column), our **FlowP-SAM+FlowI-SAM (seq)** accurately separates and segments all moving objects.

achieves outstanding performance across single- and multi-object benchmarks. Note, DAVIS17 annotations are class-based, which could be unclear and inconsistent with the class-agnostic unsupervised VOS methods. More discussion and results are provided in the arxiv version [48].

7.6 Qualitative Visualisations

In this section, example visualisations are provided across multiple datasets. We refer more visualisations to the arxiv version [48].

Fig. 4 illustrates the segmentation predictions based on only optical flow inputs. Compared to OCLR-real, our **FlowI-SAM** accurately identifies and disentangles the moving objects from the noisy backgrounds (*e.g.*, the person in the first column and the ducks in the fourth column), as well as extracts fine structures (*e.g.*, the camouflaged insect in the fifth column) from optical flow.

Fig. 5 further provides the visualisations of the RGB-based method, where the prior work (Xie et al. [47] + SAM [13]) sometimes fails to (i) identify the moving objects (*e.g.*, missing leopard in the fifth column); (ii) distinguish between multiple objects (*e.g.*, entangled object segmentation in the second and fourth columns), while our **FlowI-SAM+FlowP-SAM** (seq) incorporates RGB-based prediction with flow prompts, resulting in the accurate localisation and segmentation of moving objects.

8 Discussion

In this paper, we focus on moving object segmentation in real-world videos, by incorporating per-frame SAM with motion information (optical flow) in two ways: (i) for flow-only segmentation, we introduce **FlowI-SAM** that directly takes in optical flow as inputs; (ii) for RGB-based segmentation (**FlowP-SAM**), we utilise motion information to generate flow prompts as guidance. The former (**FlowI-SAM**) is particularly effective in scenarios with predominant motion and/or where RGB information might introduce confusion, such as in moving object detection and camouflaged object discovery. Additionally, owing to simple texture and a low cross-domain gap in optical flow, it generalises to diverse domains beyond everyday videos. On the other hand, **FlowP-SAM** focuses on common videos where both RGB and motion are informative and can be utilised to resolve ambiguities for independent objects in common motion.

Both approaches deliver state-of-the-art performance in frame-level segmentation across single- and multi-object benchmarks. Additionally, we develop a frame-wise association method that amalgamates predictions from **FlowI-SAM** and **FlowP-SAM**, achieving sequence-level segmentation predictions that outperform existing methods on DAVIS16, DAVIS17-m, YTVOS18-m, and MoCA.

The major limitation of this work is its extended running time, attributed to the computationally heavy image encoder in the vanilla SAM. However, our approach is generally applicable to other prompt-based segmentation models. With the emergence of more efficient versions of SAM, we anticipate a significant reduction in inference time.

Acknowledgments. This research is supported by the UK EPSRC CDT in AIMS (EP/S024050/1), a Clarendon Scholarship, a Royal Society Research Professorship RP\R1\191132, and the UK EPSRC Programme Grant Visual AI (EP/T028572/1).

References

1. Baker, S., Roth, S., Scharstein, D., Black, M.J., Lewis, J., Szeliski, R.: A database and evaluation methodology for optical flow. In: ICCV (2007)
2. Bideau, P., Learned-Miller, E.: It’s moving! a probabilistic model for causal motion segmentation in moving camera videos. In: ECCV (2016)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
4. Cen, J., Fang, J., Yang, C., Xie, L., Zhang, X., Shen, W., Tian, Q.: Segment any 3d gaussians. arXiv preprint arXiv:2312.00860 (2023)
5. Chen, T., Zhu, L., Ding, C., Cao, R., Zhang, S., Wang, Y., Li, Z., Sun, L., Mao, P., Zang, Y.: Sam-adapter: Adapting segment anything in underperformed scenes. In: ICCV Workshop (2023)
6. Cheng, H.K., Oh, S.W., Price, B., Schwing, A., Lee, J.Y.: Tracking anything with decoupled video segmentation. In: ICCV (2023)
7. Cheng, H.K., Schwing, A.G.: XMem: Long-term video object segmentation with an atkinson-shiffrin memory model. In: ECCV (2022)
8. Cheng, Y., Li, L., Xu, Y., Li, X., Yang, Z., Wang, W., Yang, Y.: Segment and track anything. arXiv preprint arXiv:2305.06558 (2023)
9. Cho, D., Hong, S., Kang, S., Kim, J.: Key instance selection for unsupervised video object segmentation. arXiv preprint arXiv:1906.07851 (2019)
10. Cho, S., Lee, M., Lee, S., Park, C., Kim, D., Lee, S.: Treating motion as option to reduce motion dependency in unsupervised video object segmentation. In: WACV (2023)
11. Choudhury, S., Karazija, L., Laina, I., Vedaldi, A., Rupperecht, C.: Guess What Moves: Unsupervised Video and Image Segmentation by Anticipating Motion. In: BMVC (2022)
12. Jabri, A., Owens, A., Efros, A.A.: Space-time correspondence as a contrastive random walk. In: NeurIPS (2020)
13. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: ICCV (2023)
14. Lai, Z., Lu, E., Xie, W.: Mast: A memory-augmented self-supervised tracker. In: CVPR (2020)
15. Lai, Z., Xie, W.: Self-supervised learning for video correspondence flow. In: BMVC (2019)
16. Lamdouar, H., Xie, W., Zisserman, A.: Segmenting invisible moving objects. In: BMVC (2021)
17. Lamdouar, H., Yang, C., Xie, W., Zisserman, A.: Betrayed by motion: Camouflaged object discovery via motion segmentation. In: ACCV (2020)
18. Lee, M., Cho, S., Lee, S., Park, C., Lee, S.: Unsupervised video object segmentation via prototype memory network. In: WACV (2023)
19. Li, F., Kim, T., Humayun, A., Tsai, D., Rehg, J.M.: Video segmentation by tracking many figure-ground segments. In: ICCV (2013)
20. Li, S., Seybold, B., Vorobyov, A., Fathi, A., Huang, Q., Kuo, C.C.J.: Instance embedding transfer to unsupervised video object segmentation. In: CVPR (2018)
21. Lin, H., Wu, R., Liu, S., Lu, J., Jia, J.: Video instance segmentation with a propose-reduce paradigm. In: ICCV (2021)
22. Lu, X., Wang, W., Ma, C., Shen, J., Shao, L., Porikli, F.: See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In: CVPR (2019)

23. Luiten, J., Zufikar, I.E., Leibe, B.: Unovost: Unsupervised offline video object segmentation and tracking. In: WACV (2020)
24. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* (2024)
25. Mahendran, A., Thewlis, J., Vedaldi, A.: Self-supervised segmentation by grouping optical-flow. In: ECCV (2018)
26. Meunier, E., Badoual, A., Bouthemy, P.: Em-driven unsupervised learning for efficient motion segmentation. *IEEE TPAMI* (2022)
27. Meunier, E., Bouthemy, P.: Unsupervised space-time network for temporally-consistent segmentation of multiple motions. In: CVPR (2023)
28. Miao, B., Bennamoun, M., Gao, Y., Mian, A.: Self-supervised video object segmentation by motion-aware mask propagation. In: ICME (2022)
29. Ochs, P., Malik, J., Brox, T.: Segmentation of moving objects by long term video analysis. *IEEE TPAMI* (2014)
30. Ochs, P., Brox, T.: Object segmentation in video: a hierarchical variational approach for turning point trajectories into dense regions. In: ICCV (2011)
31. Oh, S.W., Lee, J.Y., Xu, N., Kim, S.J.: Video object segmentation using space-time memory networks. In: ICCV (2019)
32. Pan, X., Li, P., Yang, Z., Zhou, H., Zhou, C., Yang, H., Zhou, J., Yang, Y.: In-n-out generative learning for dense unsupervised video segmentation. In: ACM MM (2022)
33. Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: CVPR (2016)
34. Ponimatkin, G., Samet, N., Xiao, Y., Du, Y., Marlet, R., Lepetit, V.: A simple and powerful global optimization for unsupervised video object segmentation. In: WACV (2023)
35. Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Gool, L.V.: The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675* (2017)
36. Ren, S., Luzi, F., Lahrichi, S., Kassaw, K., Collins, L.M., Bradbury, K., Malof, J.M.: Segment anything, from space? In: WACV (2024)
37. Safadoust, S., Güney, F.: Multi-object discovery by low-dimensional object motion. In: ICCV (2023)
38. Sun, Y., Chen, J., Zhang, S., Zhang, X., Chen, Q., Zhang, G., Ding, E., Wang, J., Li, Z.: Vrp-sam: Sam with visual reference prompt. In: CVPR (2024)
39. Tang, L., Xiao, H., Li, B.: Can sam segment anything? when sam meets camouflaged object detection. *arXiv preprint arXiv:2304.04709* (2023)
40. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: ECCV (2020)
41. Ventura, C., Bellver, M., Girbau, A., Salvador, A., Marques, F., Giro-i Nieto, X.: RVOS: End-to-end recurrent network for video object segmentation. In: CVPR (2019)
42. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: ECCV (2018)
43. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: CVPR (2019)
44. Wang, X., Misra, I., Zeng, Z., Girdhar, R., Darrell, T.: Videocutler: Surprisingly simple unsupervised video instance segmentation. *arXiv preprint arXiv:2308.14710* (2023)

45. Wu, J., Ji, W., Liu, Y., Fu, H., Xu, M., Xu, Y., Jin, Y.: Medical sam adapter: Adapting segment anything model for medical image segmentation. arXiv preprint arXiv:2304.12620 (2023)
46. Xie, J., Xie, W., Zisserman, A.: Segmenting moving objects via an object-centric layered representation. In: NeurIPS (2022)
47. Xie, J., Xie, W., Zisserman, A.: Appearance-based refinement for object-centric motion segmentation. arXiv:2312.11463 (2023)
48. Xie, J., Yang, C., Xie, W., Zisserman, A.: Moving object segmentation: All you need is sam (and flow). arXiv preprint arXiv:2404.12389 (2024), <https://arxiv.org/abs/2404.12389>
49. Xiong, Y., Varadarajan, B., Wu, L., Xiang, X., Xiao, F., Zhu, C., Dai, X., Wang, D., Sun, F., Iandola, F., Krishnamoorthi, R., Chandra, V.: Efficientsam: Leveraged masked image pretraining for efficient segment anything. arXiv:2312.00863 (2023)
50. Xu, N., Yang, L., Fan, Y., Yue, D., Liang, Y., Yang, J., Huang, T.: Youtube-vos: A large-scale video object segmentation benchmark. In: ECCV (2018)
51. Yang, C., Lamdouar, H., Lu, E., Zisserman, A., Xie, W.: Self-supervised video object segmentation by motion grouping. In: ICCV (2021)
52. Yang, S., Zhang, L., Qi, J., Lu, H., Wang, S., Zhang, X.: Learning motion-appearance co-attention for zero-shot video object segmentation. In: ICCV (2021)
53. Yang, Y., Lai, B., Soatto, S.: Dystab: Unsupervised object segmentation via dynamic-static bootstrapping. In: CVPR (2021)
54. Yang, Z., Wang, Q., Bertinetto, L., Bai, S., Hu, W., Torr, P.H.: Anchor diffusion for unsupervised video object segmentation. In: ICCV (2019)
55. Yang, Z., Yang, Y.: Decoupling features in hierarchical propagation for video object segmentation. In: NeurIPS (2022)
56. Zhang, C., Han, D., Qiao, Y., Kim, J.U., Bae, S.H., Lee, S., Hong, C.S.: Faster segment anything: Towards lightweight sam for mobile applications. arXiv preprint arXiv:2306.14289 (2023)
57. Zhang, K., Zhao, Z., Liu, D., Liu, Q., Liu, B.: Deep transport network for unsupervised video object segmentation. In: ICCV (2021)
58. Zhang, X., Gu, C., Zhu, S.: Sam-helps-shadow: when segment anything model meet shadow removal. arXiv preprint arXiv:2306.06113 (2023)
59. Zhang, Z., Zhang, S., Wei, Z., Dai, Z., Zhu, S.: Uvosam: A mask-free paradigm for unsupervised video object segmentation via segment anything model. arXiv preprint arXiv:2305.12659 (2024)
60. Zhao, X., Ding, W., An, Y., Du, Y., Yu, T., Li, M., Tang, M., Wang, J.: Fast segment anything. arXiv preprint arXiv:2306.12156 (2023)
61. Zheng, Z., Zhong, Y., Zhang, L., Ermon, S.: Segment any change. arXiv:2402.01188 (2024)
62. Zhou, T., Wang, S., Zhou, Y., Yao, Y., Li, J., Shao, L.: Motion-attentive transition for zero-shot video object segmentation. In: AAAI (2020)
63. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once. arXiv preprint arXiv:2304.06718 (2023)