

LiVOS: Light Video Object Segmentation with Gated Linear Matching

Qin Liu¹ Jianfeng Wang² Zhengyuan Yang² Linjie Li²
Kevin Lin² Marc Niethammer¹ Lijuan Wang²
¹UNC Chapel Hill ²Microsoft

<https://github.com/uncbiag/LiVOS>

Abstract

*Semi-supervised video object segmentation (VOS) has been largely driven by space-time memory (STM) networks, which store past frame features in a spatiotemporal memory to segment the current frame via softmax attention. However, STM networks face memory limitations due to the quadratic complexity of softmax matching, restricting their applicability as video length and resolution increase. To address this, we propose **LiVOS**, a lightweight memory network that employs linear matching via linear attention, reformulating memory matching into a recurrent process that reduces the quadratic attention matrix to a constant-size, spatiotemporal-agnostic 2D state. To enhance selectivity, we introduce gated linear matching, where a data-dependent gate matrix is multiplied with the state matrix to control what information to retain or discard. Experiments on diverse benchmarks demonstrated the effectiveness of our method. It achieved **64.8 J&F** on MOSE and **85.1 J&F** on DAVIS, surpassing all non-STM methods and narrowing the gap with STM-based approaches. For longer and higher-resolution videos, it matched STM-based methods with **53% less GPU memory** and supports **4096p** inference on a 32G consumer-grade GPU—a previously cost-prohibitive capability—opening the door for long and high-resolution video foundation models.*

1. Introduction

Video object segmentation (VOS) involves separating specific objects from the background in a sequence of video frames. It plays a critical role in various real-world applications, including robotics [41], video editing [28], and medical imaging [45]. Based on the level of user input, VOS tasks are generally categorized into three settings: unsupervised, semi-supervised (or one-shot), and interactive [20]. In this work, we focus on the semi-supervised setting, where the segmentation of the first frame is provided, and the model is required to propagate this segmentation across subsequent frames.

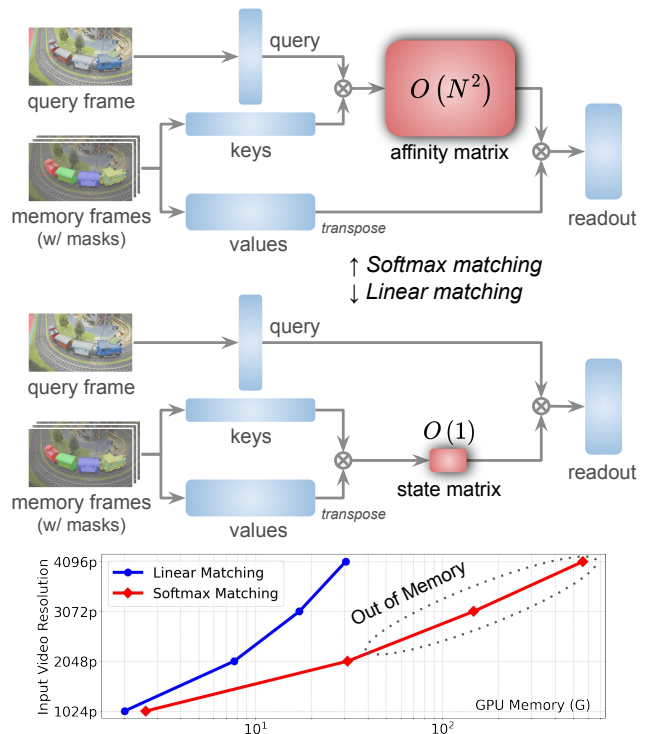


Figure 1. **Top:** Conceptual comparison of *softmax* vs. *linear* matching in video object segmentation. **Bottom:** Softmax matching suffers from memory constraints due to its *quadratic* attention complexity, while linear matching achieves linear growth with a *constant-size* state. Curves are based on results in Tab. 3.

Semi-supervised VOS is primarily driven by space-time memory (STM) networks [10, 12, 14, 36, 43], which store past frames and their segmentations as *keys* and *values* in the memory bank to help segment the *query* frame. Segmentation occurs by matching each pixel in the query frame to all spatiotemporal pixels in the memory frames via softmax attention [49]. This *softmax matching* process is highly inefficient due to the large attention matrix with space and time complexities of $\mathcal{O}(HW \times THW)$, where H and W indicate the spatial dimensions of an image and T is the num-




Figure 2. CPU latency comparison between softmax matching and linear matching. Softmax attention scales linearly over time (i.e., the number of memory frames T) and quadratically with input resolution HW . Latency is measured on an Intel Core-i7 (2.80GHz) CPU with PyTorch 2.0, batch size 1, and fp32.

ber of memory frames. Here the number of objects is not considered for simplicity. As shown in Fig. 2, the CPU latency of softmax matching increases *linearly* over time and *quadratically* with spatial dimensions. This is problematic as video length or resolution increase, making computations too slow or causing out-of-memory issues (Fig. 1).

To alleviate these issues, one could limit the number of memory frames or downsample the video’s spatial resolution. However, a fixed-size memory bank may lead to failures in challenging scenarios such as for occlusions or fast-moving objects [10]. Further, downsampling a video’s spatial resolution may lose fine details in object masks, as shown in Fig. 3, leading to inaccurate segmentations. Therefore, memory networks relying on softmax matching face significant challenges as video lengths and resolutions increase, limiting their applicability on resource-constrained devices.

To address these challenges, strategies such as knowledge distillation [34], recurrent feature embedding [29], and spatiotemporal redundancy compressing [53] have been explored. However, these methods still rely on *softmax* matching, which we identify as the *core* memory bottleneck. We address this limitation by introducing *linear* matching, reformulating the memory matching process into a recurrent framework that reduces the large attention matrix to a constant-size 2D state, yielding space and time complexities of $\mathcal{O}(THW)$. To enhance selectivity, we further introduce gated linear matching, where the state is multiplied by a data-dependent transition matrix [58] that determines what information to retain or discard. Additionally, we leverage lightweight sensory memory [10] and object memory [14] to improve performance (we claim no contribution for these components). We thus present **LiVOS**, the first light memory network that maintains *constant* memory usage for arbitrarily long videos and achieves *linear* memory growth as video resolution increases.

We evaluated LiVOS on a diverse range of video benchmarks, including both short- and long-term video datasets, as well as high-resolution videos. It achieved **64.8** $\mathcal{J}\&\mathcal{F}$

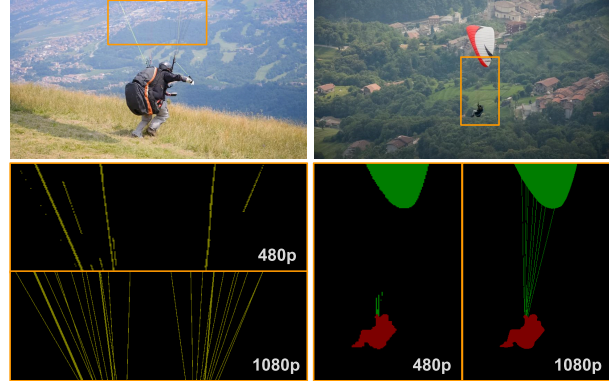


Figure 3. Masks of thin structures at different resolutions. Thin structures may lose fine details at 480p, the standard resolution for VOS. However, existing memory networks relying on softmax matching are not efficient for processing high resolution videos.

on MOSE and **85.1** $\mathcal{J}\&\mathcal{F}$ on DAVIS, surpassing all non-STM methods and narrowing the gap with STM-based approaches. For longer and higher-resolution videos, it matched STM-based methods with **53%** less GPU memory and supports **4096p** inference on a 32G consumer-grade GPU, which is not easily possible with softmax attention. While our method is not optimized for high-resolution videos, it paves the way for developing foundation models tailored to the high-resolution regime.

In summary, we introduce a lightweight memory network that: 1) employs *linear* instead of *softmax* attention for memory matching, enabling efficient video object segmentation, 2) maintains constant memory usage regardless of video length, and 3) supports 4096p high-resolution inference on 32G consumer-grade hardware. Our method demonstrates competitive performance across various video benchmarks, offering computational efficiency compared to state-of-the-art methods.

2. Related Work

Semi-supervised VOS. Semi-supervised VOS methods propagate sparse manual annotations, typically provided as one or more labeled frames, across the entire video sequence. Note that the term “semi-supervised” refers to the level of supervision required during inference, not training [39]. Early VOS methods [1, 2, 25, 33, 38] solve an optimization problem with an energy defined over a graph structure with hand-crafted energy terms. With the success of deep learning, various deep networks [3–5, 7, 24, 30, 47, 48, 50, 52–55] have been proposed, including propagation-based [40], detection-based [7], and hybrid methods [35, 57]. Among these deep networks, on-line learning [7, 32, 51] is a common practice in which networks are trained or fine-tuned during test time. However,

online learning results in slow inference, which limits its practical use. To eliminate the need for online adaptation, offline methods [10–14, 36, 59] design deep networks that can perform object-agnostic segmentation at test time. Our method belongs to offline learning approaches; specifically, it uses a memory network with linear matching.

Memory-based VOS. Space-time memory (STM) networks [10, 12, 14, 36, 61] are popular memory-based methods. The seminal work of Oh *et al.* [36] stores past frames and their object masks as key-value pairs to help segment the current frame, treated as a query. Segmentation is performed by matching each spatial position in the query to all spatiotemporal positions in the memory frames via softmax attention [49]. This softmax matching mechanism has been widely adopted in subsequent works [10, 12, 14, 43], with improved computational efficiency [12, 61] and hierarchical memory storage [10]. Recently, Cutie [14] achieved state-of-the-art performance by enhancing memory read-out with an object transformer using cross attention. However, such methods encounter high computational costs and out-of-memory issues as video length and resolution increase [10, 12–14, 29]. In contrast, our method applies linear attention for memory matching, reducing the large softmax attention matrix to a constant-size 2D state. This enables constant computational costs for arbitrarily long videos and greatly improves memory efficiency as resolution scales.

Efficient attention. Traditional softmax attention [49] suffers from significant memory bottlenecks with high-resolution videos segmentation, as it requires storing an attention matrix that is *quadratic* to the number of pixels in each frame. Linear attention [27] reduces the computational complexity of attention from *quadratic* to *linear* by linearizing the attention, which can be computed recurrently as a linear recurrent neural network. While linear attention is computationally more efficient, it often comes with a trade-off in model performance. Shen *et al.* [44] improve linear attention by proposing a separate normalization of queries and keys using a softmax function before computing the inner product. Cai *et al.* [8] propose multi-scale linear attention for high-resolution dense prediction tasks, achieving state-of-the-art performance. Flash Attention [16] accelerates GPU training by optimizing the use of memory hierarchy, while sparse-local attention [15] and sliding window attention [26] reduce global attention costs through selective focusing on parts of the spatio-temporal volume. Recently, Yang *et al.* [58] introduced gated linear attention, an improved version of linear attention with data-dependent gates, delivering competitive performance in language modeling compared to standard softmax attention and other linear-time-inference models such as Ret-

Net [46] and Mamba [21]. Building on these advances, we investigate linear attention for semi-supervised VOS to overcome the scalability challenges of softmax-based memory networks in long, high-resolution settings.

Efficient VOS. With the rise of long and high-resolution videos, efficiency in VOS has become crucial to handle larger data volumes and computational demands while ensuring high-quality performance. Early propagation-based methods [9, 57] are fast and memory-efficient but less accurate. Recent memory networks [12, 36] improve efficiency by updating the memory bank every five frames for long videos. XMem [10] enhances memory efficiency further by bounding memory bank size and introducing a training-free memory consolidation mechanism for long-term VOS, along with a sensory memory for temporal smoothness. This hierarchical memory design is adopted by Cutie [14], which uses cross-attention in an object transformer to boost performance. However, Cutie’s fixed-size memory bank still faces computational limits with increasing resolution and struggles with occlusions and fast motion. Our approach overcomes these limitations by replacing softmax matching, the primary memory bottleneck, with linear matching, ensuring constant costs for long videos and linear memory growth as resolution scales.

3. Method

We propose LiVOS, a light memory network for video object segmentation. The key innovation lies in replacing *softmax* matching with *linear* matching in standard space-time memory networks, effectively eliminating the large attention matrix that grows linearly over time and quadratically with spatial dimensions. With linear matching, LiVOS maintains a *constant-size* state matrix that is updated recurrently, making it highly efficient and well-suited for long and high-resolution videos. In the following sections, we first outline the architecture of the light memory networks (Sec. 3.1), then present gated linear matching (Sec. 3.2), an enhanced form of linear matching at the core of our method, and conclude with implementation details (Sec. 3.4).

3.1. Light Memory Networks

As discussed in Sec. 1, classical space-time memory (STM) networks are highly inefficient due to the *quadratic* complexity of softmax attention, leading to memory issues as video length and resolution grow. To address this, we reduce computational complexity by using *linear* attention for memory matching, replacing the large attention matrix with a *lightweight, constant-size* state matrix. This design forms the basis of our light memory network, as illustrated in Fig. 4. We first encode video frames into keys with an image encoder and masks into values with a mask encoder.

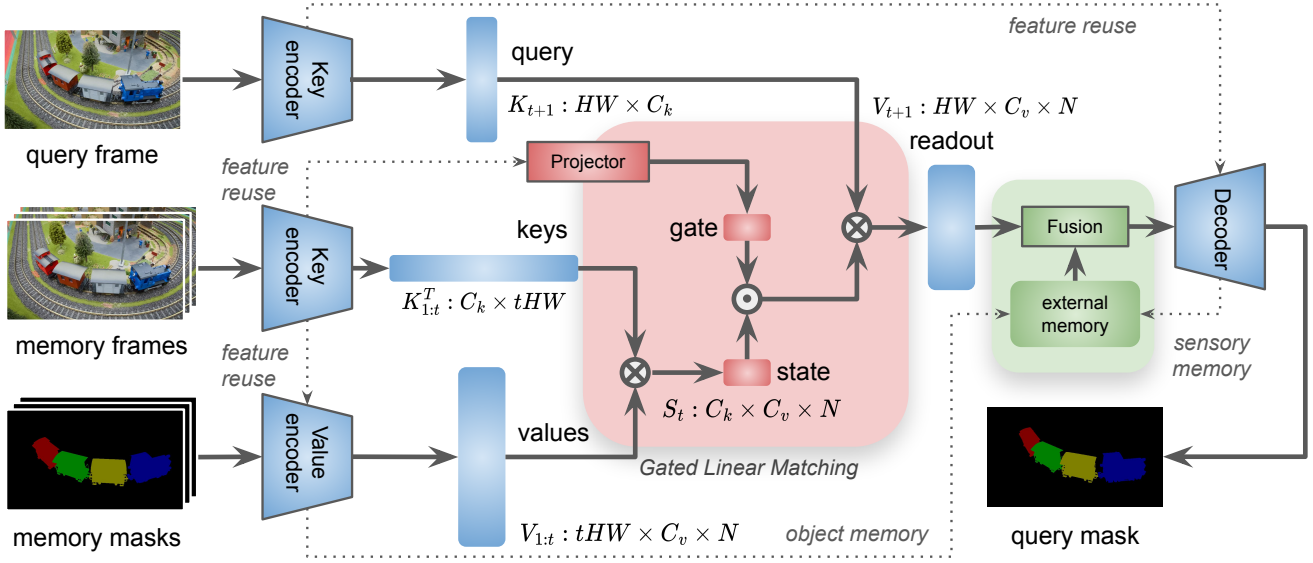


Figure 4. *LiVOS Overview*. Given a query frame, we first extract its key using an image encoder and retrieve its value via gated linear matching. This value is then enhanced by two external memories before being processed by a lightweight mask decoder for segmentation. Notably, during memory matching, our method requires only a constant-size state and gate matrix. The gate matrix is generated by a gate projector that converts the pixel-level features from the last memory frame into a sparse vector, which is then transformed to a gate matrix for element-wise multiplication with the state.

We then employ gated linear matching, an enhanced form of linear matching, to generate the query frame’s memory readout. Following Cutie [14], we enrich the memory readout by two lightweight memories: a sensory memory [10] for low-level object information and an object memory [14] for high-level object semantics. The sensory memory is fused with the readout via element-wise addition, while the object memory is fused via an object transformer consisting with cross-attention layers. The enriched readout is then processed by a mask decoder to generate the final segmentation mask. Additionally, we reuse features from the key encoder to enhance the features for both the value encoder and the mask decoder. In the next section, we detail gated linear matching, the core of our light memory network.

3.2. Gated Linear Matching

Notations. Given t memory frames $\mathbf{I}_{1:t}$ and their segmentations $\mathbf{M}_{1:t}$, we extract from them memory keys $\mathbf{K}_{1:t} \in \mathbb{R}^{tHW \times C_k}$ and values $\mathbf{V}_{1:t} \in \mathbb{R}^{tHW \times C_v}$. Here, H and W denotes the spatial dimensions, and C_k and C_v are the feature dimensions for the keys and values, respectively. Given a query frame \mathbf{I}_{t+1} , we aim to obtain its segmentation \mathbf{M}_{t+1} . We first extract its key $\mathbf{K}_{t+1} \in \mathbb{R}^{HW \times C_k}$ and then retrieve its value $\mathbf{V}_{t+1} \in \mathbb{R}^{HW \times C_v}$ through gated linear matching. To illustrate this, we begin by introducing the traditional softmax matching.

Softmax Matching. Softmax matching is widely used in space-time memory (STM) networks [10, 12, 14, 36]. Suppose we have a spatiotemporal memory consisting of memory keys $\mathbf{K}_{1:t} \in \mathbb{R}^{tHW \times C_k}$ and values $\mathbf{V}_{1:t} \in \mathbb{R}^{tHW \times C_v}$. Given a query frame \mathbf{I}_{t+1} , we first extract its key $\mathbf{K}_{t+1} \in \mathbb{R}^{HW \times C_k}$ and then retrieve its value $\mathbf{V}_{t+1} \in \mathbb{R}^{HW \times C_v}$ via softmax attention:

$$\mathbf{V}_{t+1} = \text{Softmax}(\mathbf{K}_{t+1}\mathbf{K}_{1:t}^T)\mathbf{V}_{1:t}. \quad (1)$$

We reformulate the above softmax attention in a parallel form for t memory frames:

$$\mathbf{V}_{t+1} = \frac{\sum_{i=1}^t \exp(\mathbf{K}_{t+1}\mathbf{K}_i^T)\mathbf{V}_i}{\sum_{i=1}^t \exp(\mathbf{K}_{t+1}\mathbf{K}_i^T)\mathbf{1}}, \quad (2)$$

where $\exp(\cdot)$ is an element-wise exponential for the matrix entries and $\mathbf{1} \in \mathbb{R}^{HW \times 1}$ is an all-one vector. The division in the above equation is applied element-wise over the spatial dimensions HW . Unless specified otherwise, we will continue using this element-wise division throughout the manuscript to keep the notation concise.

Based on Eq. 2, we can write a generalized memory matching for any similarity function as follows,

$$\mathbf{V}_{t+1} = \frac{\sum_{i=1}^t \text{sim}(\mathbf{K}_{t+1}, \mathbf{K}_i)\mathbf{V}_i}{\sum_{i=1}^t \text{sim}(\mathbf{K}_{t+1}, \mathbf{K}_i)\mathbf{1}}. \quad (3)$$

Eq. 3 is equivalent to Eq. 2 if we use the similarity function $\text{sim}(\mathbf{K}_{t+1}, \mathbf{K}_i) = \exp(\mathbf{K}_{t+1}\mathbf{K}_i^T)$.

Linear Matching (Parallel Form). The similarity function $\text{sim}(\cdot)$ in Eq. 3 needs to be non-negative. Given a kernel $k(x, y) = \langle \phi(x), \phi(y) \rangle$ with a feature representation $\phi(\cdot)$, we can rewrite Eq. 3 as follows,

$$\mathbf{V}_{t+1} = \frac{\sum_{i=1}^t \phi(\mathbf{K}_{t+1}) \phi(\mathbf{K}_i)^T \mathbf{V}_i}{\sum_{i=1}^t \phi(\mathbf{K}_{t+1}) \phi(\mathbf{K}_i)^T \mathbf{1}}, \quad (4)$$

and then further simplify it by making use of the associative property of matrix multiplication to

$$\mathbf{V}_{t+1} = \frac{\phi(\mathbf{K}_{t+1}) \sum_{i=1}^t \phi(\mathbf{K}_i)^T \mathbf{V}_i}{\phi(\mathbf{K}_{t+1}) \sum_{i=1}^t \phi(\mathbf{K}_i)^T \mathbf{1}}. \quad (5)$$

$\phi(\cdot)$ can be any non-linear function. Inspired by Shen *et al.* [44], we set $\phi(\cdot)$ as softmax function that is applied row-wise to the keys $\mathbf{K}_i \in \mathbb{R}^{HW \times C_k}$.

Linear Matching (Recurrent Form). Letting $\mathbf{S}_t = \sum_{i=1}^t \phi(\mathbf{K}_i)^T \mathbf{V}_i$ and $\mathbf{Z}_t = \sum_{i=1}^t \phi(\mathbf{K}_i)^T \mathbf{1}$ where $\mathbf{S}_t \in \mathbb{R}^{C_k \times C_v \times N}$ and $\mathbf{Z}_t \in \mathbb{R}^{C_k \times HW}$, we can rewrite Eq. 5 in a recurrent form,

$$\mathbf{S}_t = \mathbf{S}_{t-1} + \phi(\mathbf{K}_i)^T \mathbf{V}_i, \quad (6)$$

$$\mathbf{Z}_t = \mathbf{Z}_{t-1} + \phi(\mathbf{K}_i)^T \mathbf{1}, \quad (7)$$

$$\mathbf{V}_{t+1} = \frac{\phi(\mathbf{K}_{t+1}) \mathbf{S}_t}{\phi(\mathbf{K}_{t+1}) \mathbf{Z}_t}. \quad (8)$$

Here, $\mathbf{S}_t \in \mathbb{R}^{C_k \times C_v}$ can be regarded as a 2D recurrent hidden state that is constant size. From a sequence modeling perspective [27], $\phi(\mathbf{K}_i)^T \mathbf{V}_i$ expands the state from the key space \mathbb{R}^{C_k} and value space \mathbb{R}^{C_v} into the state space $\mathbb{R}^{C_k \times C_v}$, while $\phi(\mathbf{K}_{t+1}) \mathbf{S}_t$ reduces it back.

Complexity Analysis. For *softmax* matching in Eq. 2, the space and time complexity is $\mathcal{O}(THW \times HW \times N)$, where T is the number of memory frames, H and W are spatial dimensions, and N is the number of objects. It is *quadratic* in the spatial dimensions HW because the attention matrix must be stored to compute the weights for values. In contrast, *linear* matching has space and time complexity of $\mathcal{O}(HW \times N)$. It is *linear* in the spatial dimensions HW because we reduce the large attention matrix to a constant-size 2D state \mathbf{S}_t . Fig. 2 shows the CPU latency comparison between linear and softmax matching.

Gated Linear Matching. The linear matching in Eq. 8 does not have a selection mechanism, which has shown to be crucial in long-context tasks [6]. To enhance selectivity, we propose gated linear matching in Eq. 9, where the state $\mathbf{S}_t \in \mathbb{R}^{C_k \times C_v}$ is multiplied element-wise by a data-dependent forget gate $\mathbf{G}_t \in (0, 1)^{C_k \times C_v}$.

$$\mathbf{S}_t = \mathbf{G}_t \odot \mathbf{S}_{t-1} + \phi(\mathbf{K}_i)^T \mathbf{V}_i, \quad (9)$$

where \odot denotes element-wise multiplication. The gate \mathbf{G}_t controls which information is enhanced or suppressed (more details please refer to Sec. 3.4). Although the concept of gated linear matching is inspired by gated linear attention [58] in language modeling, we explore its application in the context of semi-supervised VOS.

3.3. Fusion with External Memory

After gated linear matching, we obtain the memory readout \mathbf{V}_{t+1} for the query frame \mathbf{I}_{t+1} . Following Cutie [14], we enrich it with two lightweight external memories: a sensory memory for low-level object information and an object memory for high-level object semantics. We present ablation studies for these two modules in Sec. 4.4. Since we claim no contributions for these components, we provide additional details in the Appendix.

3.4. Implementation Details

Encoder. We use ResNet-50 [22] as the image encoder and ResNet-18 [22] as the mask encoder, each producing multi-scale features: \mathbf{f}_4 , \mathbf{f}_8 , and the coarsest feature $\mathbf{f}_{16} \in \mathbb{R}^{HW \times C_f}$, where the subscript denotes the stride 16. Keys $\mathbf{K} \in \mathbb{R}^{HW \times C_k}$ are extract from the image encoder’s coarsest feature using convolutional layers, and values $\mathbf{V} \in \mathbb{R}^{HW \times C_v}$ are similarly extracted from the mask encoder’s coarsest feature. Following Cutie [14], the image encoder’s multi-scale features are reused in the mask encoder for value extraction, with $C_k = 64$ and $C_v = 256$.

Decoder. We adopt the lightweight mask decoder from Cutie [14] for simplicity and efficiency. This decoder takes as input the value readout \mathbf{V} at stride 16 and skip connections \mathbf{f}_4 and \mathbf{f}_8 from the image encoder at strides 4 and 8, respectively. The value readout is processed with two up-sampling blocks, incorporating skip connections to retain high-frequency details. Each block bilinearly upsamples the input feature by a factor of two, then adds the result to the skip-connection features. For multi-object cases, we apply soft-aggregation [36] to merge object logits.

Gate. The state gate can be parameterized in various ways [58], depending on a balance between parameter efficiency, training efficiency, and state size. In our work, we obtain the data-dependent gate via low-rank parameterization $\mathbf{G}_t \in \mathbb{R}^{C_k \times C_v} = \alpha_t \mathbf{1}^T$, where $\mathbf{1} \in \mathbb{R}^{C_v \times 1}$ is an all-one vector and $\alpha_t \in (0, 1)^{C_k \times 1}$ is extracted from the image encoder’s coarsest feature map $\mathbf{f}_t \in \mathbb{R}^{HW \times C_f}$. To implement this, we first use a depth-wise convolutional layer to convert the feature \mathbf{f}_t from $\mathbb{R}^{HW \times C_f}$ to $\mathbb{R}^{HW \times C_k}$, and then sum over its spatial dimensions to obtain $\hat{\alpha}_t \in \mathbb{R}^{C_k \times 1}$, followed by a Sigmoid function to obtain $\alpha_t \in (0, 1)^{C_k \times 1}$.

Pretraining. Our method relies solely on ImageNet [17] pretraining, without additional pretraining on static images or synthetic videos like BL30K [34]. While static image pretraining is standard in baseline methods (Sec. 4.2), we omit it due to its marginal improvement and increased computational burden.

Training. We use a combined loss function of cross-entropy and soft dice loss with equal weighting following [10, 13, 14]. Following Cutie [14], we adopt point supervision with $K = 12544$ sampled points instead of the whole mask for efficient training. We use PyTorch [37] and AdamW [31] optimizer with an initial learning rate of $1e^{-4}$, a batch size of 16, and a weight decay of 0.001. Each training batch contains 8 frames randomly selected from a video and cropped to a default size 480×480 . Our training lasts for 125K iterations, reducing the learning rate by a factor of 10 after 100K and 115K iterations. The state and the gate are updated for each frame in a training batch. To mitigate overfitting, we apply a learning rate multiplier $r = 0.1$ to the image encoder. We also clip the global gradient norm to $\tau = 3$ and use stable data augmentation [13]. The training processes were conducted on four NVIDIA A6000 GPUs.

Inference. Unless stated otherwise, we resize the input videos so that the shorter edge is no more than 480 pixels, then rescale the model’s prediction to the original resolution. We segment video frames sequentially and update the state as well as the gate for each frame. For datasets such as YouTube-VOS [56] where new objects may appear in intermediate frames, we create a new state for each new object with negligible computational costs. Note that training has no such issue because we fix the number of classes in a training batch. The evaluations were conducted using an NVIDIA A6000 GPU.

4. Experiments

4.1. Benchmarks

Datasets. We train and evaluate our method on the following datasets: 1) **DAVIS 2017** [39, 42], a widely used VOS dataset with three main sets: 1) a training set with 60 videos (4,219 frames, 138 objects), 2) a validation set with 30 videos (2,023 frames, 59 objects), and 3) a test set with 30 videos (2,037 frames, 89 objects). 2) **YouTube-VOS 2019** [56], one of the largest VOS dataset, featuring a wide range of challenging real-world scenarios. The dataset consists of a training set with 3,471 videos (197,292 frames) and a validation set with 507 videos (19,981 frames). There are 65 seen object categories in the training set and 26 unseen categories only appear in the validation set, challenging a model’s generalization ability. 3) **MOSE** [18], a large-scale video object segmentation (VOS) benchmark empha-

sizing multi-object scenarios in diverse environments. It comprises 2,149 videos, divided into 1,507 training videos, 311 validation videos, and 331 testing videos. The dataset includes 5,200 objects across 36 distinct categories. In this work, we utilize the training set for model training and evaluate performance on the validation set. 4) **LVOS** [23], a dataset designed specifically for video object segmentation in long, continuous video sequences, focusing on more extended temporal challenges. It’s divided into three main sets: 1) a training set with 120 videos, 2) a validation set with 50 videos, and 3) a test set with 50 videos. There are 126,280 frames and 156,432 annotations in total. In this work, we report evaluation results on the validation and test sets without training our model on LVOS [23].

Evaluation Metrics. We report Region Jaccard \mathcal{J} [39], contour accuracy \mathcal{F} [19], and their average $\mathcal{J}\&\mathcal{F}$ to measure the segmentation quality. For YouTube-VOS, we additionally report \mathcal{G} as the average of $\mathcal{J}\&\mathcal{F}$ across both seen and unseen classes. We also report frames per second (FPS) and GPU memory consumption (*i.e.* maximum GPU memory allocated by PyTorch [37]) to measure speed and resource usage, respectively. Both metrics are recorded on the same machine under similar conditions for all methods.

4.2. Baselines

We compare LiVOS with various types of approaches, including state-of-the-art space-time memory (STM) networks, represented by Cutie [14] and methods without using STM, represented by RDE [29]. This section only introduces several representative methods. A complete introduction is included in the appendix.

RDE [29] introduces a recurrent dynamic embedding to maintain a constant-size memory, with image and mask encoders based on ResNet-50. We use the version without BL30K [11] pretraining as a baseline.

STCN [12] is a state-of-the-art memory network building on STM [36], enhancing robustness and efficiency. We use the model without BL30K [11] pretraining.

XMem [10] targets long-term videos with a compact long-term memory, a fast-updating sensory memory, and an STCN-based working memory. As with STCN, we use the model without BL30K[11] pretraining.

Cutie [14] enhances XMem[10] with a top-down, object-level memory reading mechanism for improved video segmentation. We use its small and base models as baselines. We also limit Cutie to using a single memory frame during inference for fair comparison with non-STM methods.

4.3. Comparisons

We compare our method with baselines on three types of video benchmarks: 1) short videos, 2) long-term videos, and 3) high-resolution videos.

Method	with STM	MOSE			DAVIS-17 val			DAVIS-17 test			YouTubeVOS-2019 val				
		$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	\mathcal{G}	\mathcal{J}_s	\mathcal{F}_s	\mathcal{J}_u	\mathcal{F}_u
♫ STCN [12] <small>NeurIPS'21</small>	✓	52.5	48.5	56.6	85.4	82.2	88.6	76.1	72.7	79.6	82.7	81.1	85.4	78.2	85.9
♫ AOT [61] <small>NeurIPS'21</small>	✓	58.4	54.3	62.6	84.9	82.3	87.5	79.6	75.9	83.3	85.3	83.9	88.8	79.9	88.5
♫ XMem [10] <small>ECCV'22</small>	✓	56.3	52.1	60.6	86.2	82.9	89.5	81.0	77.4	84.5	85.5	84.3	88.6	80.3	88.6
♫ DeAOT [59] <small>NeurIPS'22</small>	✓	59.0	54.6	63.4	85.2	82.2	88.2	80.7	76.9	84.5	85.6	84.2	89.2	80.2	88.8
♫ DEVA [13] <small>ICCV'23</small>	✓	60.0	55.8	64.3	86.8	83.6	90.0	82.3	78.7	85.9	85.5	85.0	89.4	79.7	88.0
♫ Cutie-small [14] <small>CVPR'24</small>	✓	62.2	58.2	66.2	87.2	84.3	90.1	84.1	80.5	87.6	86.2	85.3	89.6	80.9	89.0
♫ Cutie-small [14] <small>CVPR'24</small>	✓	67.4	63.1	71.7	86.5	83.5	89.5	83.8	80.2	87.5	86.3	85.2	89.7	81.1	89.2
♫ Cutie-base [14] <small>CVPR'24</small>	✓	64.0	60.0	67.9	88.8	85.4	92.3	84.2	80.6	87.7	86.1	85.5	90.0	80.6	88.3
♫ Cutie-base [14] <small>CVPR'24</small>	✓	68.3	64.2	72.3	88.8	85.6	91.9	85.3	81.4	89.3	86.5	85.4	90.0	81.3	89.3
♫ CFBI [60] <small>ECCV'20</small>	✗	-	-	-	81.9	79.1	84.6	74.8	71.1	78.5	-	-	-	-	-
♫ CFBI+ [62] <small>TPAMI'21</small>	✗	-	-	-	82.9	80.1	85.7	75.6	71.6	79.6	-	-	-	-	-
♫ SwiftNet [53] <small>CVPR'21</small>	✗	-	-	-	81.1	78.3	83.9	-	-	-	77.8	77.8	81.8	72.3	79.5
♫ RDE [29] <small>CVPR'22</small>	✗	46.8	42.4	51.3	84.2	80.8	87.5	77.4	73.6	81.2	81.9	81.1	85.5	76.2	84.8
♫ MobileVOS [34] <small>CVPR'23</small>	✗	-	-	-	83.7	80.2	87.1	-	-	-	82.3	81.6	86.0	76.3	85.2
♫ Cutie-small [†] [14] <small>CVPR'24</small>	✗	49.3	45.4	53.2	76.4	73.0	79.8	71.6	67.9	75.3	79.0	78.0	81.9	74.5	81.8
♫ Cutie-small [†] [14] <small>CVPR'24</small>	✗	51.7	47.7	55.7	74.9	71.7	78.2	72.3	68.8	75.9	78.9	77.6	81.4	74.6	81.9
♫ Cutie-base [†] [14] <small>CVPR'24</small>	✗	50.6	46.6	54.6	79.3	75.8	82.7	73.5	70.0	77.0	80.1	79.1	83.3	75.2	82.7
♫ Cutie-base [†] [14] <small>CVPR'24</small>	✗	52.6	48.6	56.6	77.5	74.1	80.8	73.2	69.4	76.9	79.5	78.4	82.4	75.1	82.4
♫ LiVOS	✗	59.2	55.6	62.8	84.4	81.2	87.6	78.2	74.8	81.7	79.9	82.6	86.8	71.7	78.4
♫ LiVOS*	✗	58.4	54.7	62.0	85.1	81.9	88.3	81.0	77.4	84.5	81.3	83.6	87.6	73.3	80.6
♫ LiVOS	✗	64.8	60.9	68.7	84.0	80.6	87.3	79.6	76.2	82.9	82.6	83.5	87.6	75.8	83.3

Table 1. *Quantitative comparisons on short videos (480p)*. All methods are trained on 480p videos and use ResNet as the backbone: Cutie-small uses ResNet-18, CFBI/CFBI+ use ResNet-101, and the rest use ResNet-50. ♫ represents models trained on YouTube VOS and DAVIS; ♫ denotes additional training with MOSE; * denotes training with 30k additional iterations beyond the standard 125k; † indicates downgraded Cutie models using only one memory frame during inference (more comparisons with the downgraded Cutie models are shown in the appendix). LiVOS notably improves segmentation quality over non-STM approaches and matches STM-based methods.

Method	with STM	Key Encoder	Value Encoder	LVOS val					LVOS test				
				$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	Mem↓	FPS↑	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	Mem↓	FPS↑
♫ DEVA [14] <small>ICCV'23</small>	✓	RN-50	RN-18	58.3	52.8	63.8	1147M	48.3	54.0	49.0	59.0	1347M	46.6
♫ DEVA [14] <small>ICCV'23</small>	✓	RN-50	RN-18	55.9	51.1	60.7	1147M	48.3	56.5	52.2	60.8	1347M	46.6
♫ Cutie-small [14] <small>CVPR'24</small>	✓	RN-18	RN-18	58.8	54.6	62.9	1013M	34.0	57.2	53.7	60.7	1094M	32.8
♫ Cutie-small [14] <small>CVPR'24</small>	✓	RN-18	RN-18	60.7	55.6	65.8	1013M	34.0	56.9	53.5	60.2	1094M	32.8
♫ Cutie-base [14] <small>CVPR'24</small>	✓	RN-50	RN-18	60.1	55.9	64.2	1092M	30.1	56.2	51.8	60.5	1175M	29.7
♫ Cutie-base [14] <small>CVPR'24</small>	✓	RN-50	RN-18	63.5	59.1	67.9	1092M	30.1	63.6	59.1	68.0	1175M	29.7
♫ RDE [29] <small>CVPR'22</small>	✗	RN-50	RN-18	47.2	41.7	52.7	9.0G [‡]	40.6	44.7	39.2	50.2	12.7G [‡]	39.8
♫ Cutie-small [†] [14] <small>CVPR'24</small>	✗	RN-18	RN-18	48.2	43.9	52.5	585M	56.1	45.2	41.3	49.0	686M	55.2
♫ Cutie-small [†] [14] <small>CVPR'24</small>	✗	RN-18	RN-18	46.1	42.1	50.1	585M	56.1	44.0	40.1	48.0	686M	55.2
♫ Cutie-base [†] [14] <small>CVPR'24</small>	✗	RN-50	RN-18	49.1	44.7	53.4	668M	45.5	48.7	44.6	52.7	748M	43.0
♫ Cutie-base [†] [14] <small>CVPR'24</small>	✗	RN-50	RN-18	46.8	43.1	50.6	668M	45.5	45.6	41.6	49.7	748M	43.0
♫ LiVOS	✗	RN-50	RN-18	50.6	46.5	54.7	503M	47.3	44.6	41.2	47.9	575M	45.2
♫ LiVOS*	✗	RN-50	RN-18	51.2	46.8	55.6	503M	47.3	50.9	47.0	54.7	575M	45.2
♫ LiVOS	✗	RN-50	RN-18	51.2	47.3	55.1	503M	47.3	47.0	44.0	50.0	575M	45.2

Table 2. *Quantitative comparisons on long-term videos (480p)*. We present the same model variants and notations as in Tab. 1. We additionally report FPS and the maximum GPU memory usage. ‡unexpectedly large GPU memory consumption for RDE, likely due to implementation issues in its released code. LiVOS consistently achieves the best performance across non-STM methods on both LVOS validation and test sets while using less GPU memory.

Short Videos. Tab. 1 shows results on standard VOS benchmarks: MOSE [18], DAVIS [42], and YouTube VOS [56]. Our method achieves a score of 64.8 $\mathcal{J}\&\mathcal{F}$ on MOSE [18], trailing the state-of-the-art method Cutie [14] by just 3.5 $\mathcal{J}\&\mathcal{F}$. Note that MOSE is very challenging because of the complexity of multi-object interactions and dynamic

environments. By default, Cutie [14] uses five memory frames. When limited to one memory frame, in a variant labeled Cutie[†], its performance drops significantly, falling below ours. Moreover, our method consistently outperforms other non-STM methods, including MobileVOS [34] and RDE [29], across most benchmarks.

Method	with STM	1024P			2048P			3072P			4096P		
		Mem↓	FPS↑	$\mathcal{J}\&\mathcal{F}$	Mem↓	FPS↑	$\mathcal{J}\&\mathcal{F}$	Mem↓	FPS↑	$\mathcal{J}\&\mathcal{F}$	Mem↓	FPS↑	$\mathcal{J}\&\mathcal{F}$
♣STCN [12] <small>NeurIPS'21</small>	✓	22.6G	3.7	83.1	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣AOT [61] <small>NeurIPS'21</small>	✓	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣XMem [10] <small>ECCV'22</small>	✓	12.5G	7.3	86.8	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣DeAOT [59] <small>NeurIPS'22</small>	✓	13.0G	4.5	86.9	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣DEVA [13] <small>ICCV'23</small>	✓	15.7G	5.1	90.0	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣Cutie-small [14] <small>CVPR'24</small>	✓	9.7G	7.6	89.3	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣Cutie-base [14] <small>CVPR'24</small>	✓	9.9G	6.6	89.5	150G [‡]	<i>N/A</i>	<i>N/A</i>	735G [‡]	<i>N/A</i>	<i>N/A</i>	2815G [‡]	<i>N/A</i>	<i>N/A</i>
♣RDE [29] <small>CVPR'22</small>	✗	14.4G	10.6	76.3	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣Cutie-small [†] [14] <small>CVPR'24</small>	✗	2.3G	11.9	83.7	30.3G	2.1	79.8	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>	<i>OOM</i>	<i>N/A</i>	<i>N/A</i>
♣Cutie-base [†] [14] <small>CVPR'24</small>	✗	2.6G	10.2	86.0	31.1G	1.9	79.6	147G [‡]	<i>N/A</i>	<i>N/A</i>	563G [‡]	<i>N/A</i>	<i>N/A</i>
♣LiVOS	✗	2.0G	12.8	85.0	7.7G	3.5	80.0	17.2G	1.5	73.4	30.4G	0.8	61.5

Table 3. *Quantitative comparisons on high-resolution videos ($\geq 1024p$).* These videos are upsampled from DAVIS 2017 validation set (480p). All models were tested with an NVIDIA A6000 GPU. *OOM* indicates out-of-memory, exceeding the A6000’s 48G limit. [‡] denotes estimated memory usage. Note that LiVOS utilizes no modules dedicated for high-res segmentation, leaving room for future improvements.

Long-term Videos. Tab. 2 shows results on the long video benchmark LVOS [23]. Compared with the state-of-the-art recurrent memory network RDE [29], our method is better across all benchmarks. Compared with state-of-the-art STM-based methods, our method delivers competitive performance while being faster and more memory efficient. Even with Cutie’s memory bank limited to one memory frame, our method requires less memory while achieving better segmentation quality. Note that memory networks such as Cutie and DEVA employ a bounded memory bank (by default, the number is five), so they have no memory issues when video length increases. To highlight our memory efficiency, we also compare our method with baselines on high-resolution videos, introduced below.

High-resolution Videos. Tab. 3 shows the comparisons on high-resolution videos upsampled from the DAVIS dataset. All space-time memory networks encounter out-of-memory issues at 2048p. For example, Cutie-base requires approximately 150GB of memory to run inference at 2048p, far exceeding the capabilities of consumer-grade GPUs. In contrast, our method requires only 7.7GB for 2048p inference and is the only one capable of inference at 4096p. While not specifically optimized for high-resolution videos, our method sets the stage for developing foundation models suited to high-resolution regimes. More comparisons with downsampling baselines are shown in the appendix.

4.4. Ablations

Tab. 4 shows the results of an ablation study to validate our design choices. The models are trained on YouTube-VOS and DAVIS and evaluated on the DAVIS validation set. *No gate* is a model variant that disables the gate mechanism in linear matching. *No sensory memory* is a model variant that disables the sensory memory. *No object memory* is a model variant that disables the object memory. The gate mechanism in linear matching improves 1.4% $\mathcal{J}\&\mathcal{F}$, with only

Method	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	FPS ↑	Mem ↓
No gate	83.0	80.0	86.0	40.8	573M
No object memory	80.4	77.4	83.4	70.5	565M
No sensory memory	79.1	76.3	81.9	52.6	538M
Full	84.4	81.2	87.6	40.3	574M

Table 4. Ablation study for LiVOS on DAVIS 2017 val [39].

a minor trade-off in speed and memory usage. The integration of sensory and object memories substantially boost performance, though it comes at the cost of reduced speed and higher memory consumption. In this ablation, we retain linear matching instead of swapping it with softmax matching, as that would make our method similar to Cutie [14].

5. Limitations

Our method uses a single recurrent state for the challenging VOS task, which works well for short, low-resolution videos but may lead to suboptimal results on longer, high-resolution ones. We believe a more advanced state [21], as well as multi-scale linear attention [8], could improve performance in these cases. We leave these extensions for future work.

6. Conclusion

We proposed the first light memory network that employs gated linear matching, instead of softmax matching, for memory efficient video object segmentation. Our method reformulates the memory matching process into a recurrent framework where the large, quadratic attention matrix is reduced to a small, constant-size recurrent state. Our evaluation results on long and high-resolution videos demonstrated the effectiveness and efficiency of our method, opening the door for high-resolution video foundation models.

Acknowledgement. We thank our reviewers and ACs for their insightful suggestions. This research was, in part, funded by the National Institutes of Health (NIH) under other transactions 1OT2OD038045-01 and NIAMS 1R01AR082684. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies, either expressed or implied, of the NIH.

References

- [1] S Avinash Ramakanth and R Venkatesh Babu. Seamseg: Video object segmentation using patch seams. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 376–383, 2014. [2](#)
- [2] Vijay Badrinarayanan, Fabio Galasso, and Roberto Cipolla. Label propagation in video sequences. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3265–3272. IEEE, 2010. [2](#)
- [3] Gedas Bertasius and Lorenzo Torresani. Classifying, segmenting, and tracking object instances in video with mask propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9739–9748, 2020. [2](#)
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, page 4, 2021.
- [5] Goutam Bhat, Felix Järemo Lawin, Martin Danelljan, Andreas Robinson, Michael Felsberg, Luc Van Gool, and Radu Timofte. Learning what to learn for video object segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 777–794. Springer, 2020. [2](#)
- [6] Jacob Buckman and Carles Gelada. Linear Transformers Are Faster After All, 2024. [5](#)
- [7] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 221–230, 2017. [2](#)
- [8] Han Cai, Junyan Li, Muyan Hu, Chuang Gan, and Song Han. Efficientvit: Multi-scale linear attention for high-resolution dense prediction. *arXiv preprint arXiv:2205.14756*, 2022. [3](#), [8](#)
- [9] Xi Chen, Zuoxin Li, Ye Yuan, Gang Yu, Jianxin Shen, and Donglian Qi. State-aware tracker for real-time video object segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9384–9393, 2020. [3](#)
- [10] Ho Kei Cheng and Alexander G Schwing. XMMem: Long-term video object segmentation with an Atkinson-Shiffrin memory model. In *ECCV*, pages 640–658, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [11] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. In *CVPR*, pages 5559–5568, 2021. [6](#)
- [12] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *NeurIPS*, 34:11781–11794, 2021. [1](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [13] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Alexander Schwing, and Joon-Young Lee. Tracking anything with decoupled video segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1316–1326, 2023. [6](#), [7](#), [8](#)
- [14] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3151–3161, 2024. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [15] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019. [3](#)
- [16] Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022. [3](#)
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [6](#)
- [18] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, Philip HS Torr, and Song Bai. Mose: A new dataset for video object segmentation in complex scenes. *arXiv preprint arXiv:2302.01872*, 2023. [6](#), [7](#)
- [19] Eduardo Fernandez-Moral, Renato Martins, Denis Wolf, and Patrick Rives. A new metric for evaluating semantic segmentation: leveraging global and contour accuracy. In *2018 IEEE intelligent vehicles symposium (iv)*, pages 1051–1056. IEEE, 2018. [6](#)
- [20] Mingqi Gao, Feng Zheng, James JQ Yu, Caifeng Shan, Guiguang Ding, and Jungong Han. Deep learning for video object segmentation: a review. *Artificial Intelligence Review*, 56(1):457–531, 2023. [1](#)
- [21] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. [3](#), [8](#)
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [5](#)
- [23] Lingyi Hong, Wenchao Chen, Zhongying Liu, Wei Zhang, Pinxue Guo, Zhaoyu Chen, and Wenqiang Zhang. Lvos: A benchmark for long-term video object segmentation. *arXiv preprint arXiv:2211.10181*, 2022. [6](#), [8](#)
- [24] Yuan-Ting Hu, Jia-Bin Huang, and Alexander G Schwing. Videomatch: Matching based video object segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 54–70, 2018. [2](#)
- [25] Suyog Dutt Jain and Kristen Grauman. Supervoxel-consistent foreground propagation in video. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 656–671. Springer, 2014. [2](#)
- [26] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023. [3](#)

- [27] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020. [3](#), [5](#)
- [28] Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Occlusion-aware video object inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14468–14478, 2021. [1](#)
- [29] Mingxing Li, Li Hu, Zhiwei Xiong, Bang Zhang, Pan Pan, and Dong Liu. Recurrent dynamic embedding for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1332–1341, 2022. [2](#), [3](#), [6](#), [7](#), [8](#)
- [30] Yu Li, Zhuoran Shen, and Ying Shan. Fast video object segmentation using the global context module. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 735–750. Springer, 2020. [2](#)
- [31] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [6](#)
- [32] K-K Maninis, Sergi Caelles, Yuhua Chen, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. Video object segmentation without temporal information. *IEEE transactions on pattern analysis and machine intelligence*, 41(6):1515–1530, 2018. [2](#)
- [33] Nicolas Märki, Federico Perazzi, Oliver Wang, and Alexander Sorkine-Hornung. Bilateral space video segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 743–751, 2016. [2](#)
- [34] Roy Miles, Mehmet Kerim Yucel, Bruno Manganelli, and Albert Saà-Garriga. Mobilevos: Real-time video object segmentation contrastive learning meets knowledge distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10480–10490, 2023. [2](#), [6](#), [7](#)
- [35] Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, and Seon Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7376–7385, 2018. [2](#)
- [36] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *CVPR*, pages 9226–9235, 2019. [1](#), [3](#), [4](#), [5](#), [6](#)
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [6](#)
- [38] Federico Perazzi, Oliver Wang, Markus Gross, and Alexander Sorkine-Hornung. Fully connected object proposals for video segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 3227–3234, 2015. [2](#)
- [39] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. [2](#), [6](#), [8](#)
- [40] Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, and Alexander Sorkine-Hornung. Learning video object segmentation from static images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2663–2672, 2017. [2](#)
- [41] Vladimír Petřík, Mohammad Nomaan Qureshi, Josef Sivic, and Makar Tapaswi. Learning object manipulation skills from video via approximate differentiable physics. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7375–7382. IEEE, 2022. [1](#)
- [42] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. [6](#), [7](#)
- [43] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. [1](#), [3](#)
- [44] Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3531–3539, 2021. [3](#), [5](#)
- [45] Alexey A Shvets, Alexander Rakhlin, Alexandr A Kalinin, and Vladimir I Iglovikov. Automatic instrument segmentation in robot-assisted surgery using deep learning. In *2018 17th IEEE international conference on machine learning and applications (ICMLA)*, pages 624–628. IEEE, 2018. [1](#)
- [46] Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. Retentive network: A successor to transformer for large language models. *arXiv preprint arXiv:2307.08621*, 2023. [3](#)
- [47] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning video object segmentation with visual memory. In *Proceedings of the IEEE international conference on computer vision*, pages 4481–4490, 2017. [2](#)
- [48] Pavel Tokmakov, Jie Li, and Adrien Gaidon. Breaking the” object” in video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22836–22845, 2023. [2](#)
- [49] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. [1](#), [3](#)
- [50] Carles Ventura, Miriam Bellver, Andreu Girbau, Amaia Salvador, Ferran Marques, and Xavier Giro-i Nieto. Rvos: End-to-end recurrent network for video object segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5277–5286, 2019. [2](#)
- [51] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for video object segmentation. *arXiv preprint arXiv:1706.09364*, 2017. [2](#)
- [52] Paul Voigtlaender, Yuning Chai, Florian Schroff, Hartwig Adam, Bastian Leibe, and Liang-Chieh Chen. Feelvos:

- Fast end-to-end embedding learning for video object segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9481–9490, 2019. [2](#)
- [53] Haochen Wang, Xiaolong Jiang, Haibing Ren, Yao Hu, and Song Bai. Swiftnet: Real-time video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1296–1305, 2021. [2](#), [7](#)
- [54] Junke Wang, Dongdong Chen, Zuxuan Wu, Chong Luo, Chuanxin Tang, Xiyang Dai, Yucheng Zhao, Yujia Xie, Lu Yuan, and Yu-Gang Jiang. Look before you match: Instance understanding matters in video object segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2268–2278, 2023.
- [55] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. Learning to detect salient objects with image-level supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 136–145, 2017. [2](#)
- [56] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. [6](#), [7](#)
- [57] Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, and Aggelos K Katsaggelos. Efficient video object segmentation via network modulation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6499–6507, 2018. [2](#), [3](#)
- [58] Songlin Yang, Bailin Wang, Yikang Shen, Rameswar Panda, and Yoon Kim. Gated linear attention transformers with hardware-efficient training. *arXiv preprint arXiv:2312.06635*, 2023. [2](#), [3](#), [5](#)
- [59] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. *Advances in Neural Information Processing Systems*, 35:36324–36336, 2022. [3](#), [7](#), [8](#)
- [60] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *European Conference on Computer Vision*, pages 332–348. Springer, 2020. [7](#)
- [61] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34:2491–2502, 2021. [3](#), [7](#), [8](#)
- [62] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by multi-scale foreground-background integration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4701–4712, 2021. [7](#)