RMem: Restricted Memory Banks Improve Video Object Segmentation

Junbao Zhou*   Ziqi Pang*   Yu-Xiong Wang

University of Illinois Urbana-Champaign

{junbaoz,ziqip2,yxw}@illinois.edu

Abstract

With recent video object segmentation (VOS) benchmarks evolving to challenging scenarios, we revisit a simple but overlooked strategy: restricting the size of memory banks. This diverges from the prevalent practice of expanding memory banks to accommodate extensive historical information. Our specially designed "memory deciphering" study offers a pivotal insight underpinning such a strategy: expanding memory banks, while seemingly beneficial, actually increases the difficulty for VOS modules to decode relevant features due to the confusion from redundant information. By restricting memory banks to a limited number of essential frames, we achieve a notable improvement in VOS accuracy. This process balances the importance and freshness of frames to maintain an informative memory bank within a bounded capacity. Additionally, restricted memory banks reduce the training-inference discrepancy in memory lengths compared with continuous expansion. This fosters new opportunities in temporal reasoning and enables us to introduce the previously overlooked "temporal positional embedding." Finally, our insights are embodied in "RMem" ("R" for restricted), a simple yet effective VOS modification that excels at challenging VOS scenarios and establishes new state of the art for object state changes (on the VOST dataset) and long videos (on the Long Videos dataset). Our code and demo are available at https://restricted-memory.github.io/.

1. Introduction

The rapid progress of video object segmentation (VOS) algorithms has motivated the creation of more challenging benchmarks, as exemplified by VOST [39] on more complicated videos with significant *object state changes* and the Long Videos dataset [27] featuring extremely *long* duration. These benchmarks elevate the spatio-temporal modeling and prompt us to reassess conventional VOS designs: *can learning-based VOS modules effectively decipher historical information in such challenging scenarios?*

To delve into this issue, it is essential to focus on *memory banks*, which are central to storing past features and feeding input to VOS modules, and are fundamental in the memory-
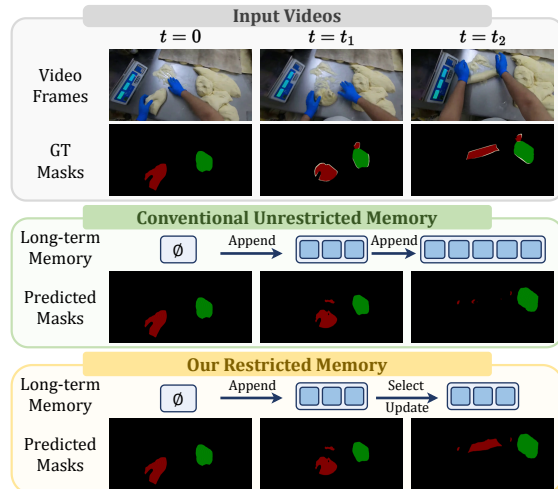
*Equal contribution.

Figure 1. In light of challenging object state changes [39, 46, 52], we rethink the conventional VOS approach of continuously accumulating the features into memory banks: despite capturing all the information, it complicates the deciphering of relevant features. Conversely, restricted memory banks significantly enhance VOS.

based VOS framework [9, 11, 50]. Typically, the memory banks are managed via the simple intuition of *expansion*, continuously *appending* newly sampled frames as the video progresses. While this approach is intended to encompass all historical information, thereby enhancing VOS, we realize its potential limitation: as videos become longer or more complex, these expanding memory banks may overwhelm the capability of VOS modules to discern reliable features.

We investigate this hypothesis by conducting a pilot study, named "*memory deciphering*," to quantify the decoding capability of VOS modules. In our analysis, we continue to use object segmentation as the proxy to VOS, but shift the prediction target to decoding *the object mask at the initial frame (frame 0)* from the memory bank. This choice is deliberate based on the principle of controlling variables: (1) In the VOS framework, the information of frame 0 is implicitly propagated to subsequent frames, ensuring the presence of relevant information for decoding; (2) This prediction target is consistent across frames and allows for a fair comparison of decoding efficacy under varying memory sizes. Intuitively, the later frames have *rigorously richer* information than the earlier frames because of

a larger memory bank, and are thus expected to produce better decoding results. However, our observation shows the opposite: *the effectiveness of VOS modules in deciphering information diminishes with increasingly large memory banks.* Intriguingly, this degradation can be mitigated by selecting a small number of relevant frames in the memory bank, and we observe a significantly better concentration of attention scores on relevant frames and regions. Therefore, our systematic study reveals a pivotal insight: *the expansion of memory banks complicates the deciphering of VOS modules* primarily due to *redundant information*.

Inspired by such an insight, we validate its practical significance through a simple approach: *restricting memory banks to a fixed number of frames.* Our concise memory bank facilitates better spatio-temporal modeling and adaptation to object transformation according to the analysis of complex object state changes [39], as illustrated in Fig. 1. The effectiveness of our method stems from a curated memory concisely focusing the attention of VOS modules on relevant information. Based on this, we delve into the updating process when new features arrive. Our strategy balances the relevance and freshness of frame features, drawing inspiration from the upper confidence bound (UCB) algorithm [3] from multi-arm bandit problems.

In addition to enhancing the accuracy, restricted memory banks reduce *discrepancies in memory lengths* between training and inference when compared with conventional methods. Typically, VOS modules are trained on short clips with a few memory frames, so our restricted memory bank better aligns with this setup, even when handling significantly longer videos during inference. This alignment opens up opportunities to revisit techniques relying on temporal synchronization between training and inference. As a compelling example, we introduce *temporal positional embedding* to explicitly capture the ordering of memory features – a critical aspect often overlooked by previous methods – leading to superior temporal reasoning.

In conclusion, we make the following contributions:

1. We introduce the novel *memory deciphering* analysis to systematically reveal the drawbacks of expanding memory banks for VOS modules in decoding information.
2. Our revisit of *restricting memory banks* notably enhances VOS accuracy for challenging cases, cooperated with a memory update strategy balancing the relevance and freshness of frames.
3. Benefiting from smaller training-inference gaps, we introduce the previously overlooked *temporal positional embedding* to capture the order of memory frames.

Collectively, our insights lead to a simple yet strong VOS method: "RMem," which is *plug-and-play* for memory-based VOS methods. Our extensive experiments show its strengths and establish new state of the art on VOST [39] for object state changes and the Long Videos dataset [27].

2. Related Work

VOS benchmarks. VOS has evolved through several benchmarks. DAVIS [35, 36] is the first exhibiting diversity and quality, surpassing early benchmarks [5, 25, 40]. YoutubeVOS [45] further scales up by collecting more videos. Although they have enabled great progress in VOS, their limited difficulty and video lengths have spurred more challenging datasets. For example, the average duration in LVOS [20] is more than 500 frames and the Long Videos dataset [27] further extends it to over 1,000 frames, and MOSE [15] increases the difficulty by selecting videos with crowds and occlusions. To evaluate our insight on the most demanding scenarios, we highlight *object state changes* involving noticeable transformations in the existence, appearance, and shapes. Studies on state changes, *e.g.*, VS-COS [52], mostly utilize ego-centric datasets [13, 14, 18]. In this paper, we primarily select the recent VOST [39]. It combines multiple datasets and provides accurate annotations. Notably, VOST shows higher complexity and longer duration than previous YoutubeVOS and DAVIS. We mainly concentrate on the challenging benchmarks.

Memory-based VOS. Memory banks are fundamental for VOS. Earlier approaches [4, 6, 30, 37, 42] treat VOS as online learning and finetune networks with memorized features. Some others [7, 21, 43, 47, 49, 51] approach VOS as template matching but struggle with occluded or dynamically changing objects. Consequently, recent methods mostly focus on memory reading via either pixel-level or object-level attention [41]. Object-level memory reading [1, 2, 12], inspired by Mask2Former [8], excels at efficiency. However, it is less effective for delicate masks or complex scenarios, *e.g.*, VOST [39], where the objects are frequently small or cluttered. In comparison, pixel-wise memory reading [9, 11, 17, 27, 33, 38, 44, 48, 50] is more adopted for its reliable segmentation and it typically associates the current frame to memory features with attention. Our work differs from previous studies by focusing more on the general insights of *drawbacks of expanding memory banks* and plug-and-play strategies to mitigate such issues, instead of dedicated memory reading architectures.

Restricted Memory Banks in VOS. Previous studies approach restricting memory banks mostly from the efficiency aspect [9, 26, 27]. A notable representative, XMem [9], adopts a hierarchical architecture with customized modifications like similarity computation and memory potentiation. In contrast to these prior efforts, our work stands out by explicitly revealing and highlighting the *accuracy* benefits of restricted memory banks through reducing redundant information, rather than emphasizing *efficiency*. Moreover, our RMem demonstrates such an insight with a *simple plug-and-play* enhancement to the VOS framework, avoiding any noticeable increase or reliance on special operators as in XMem.
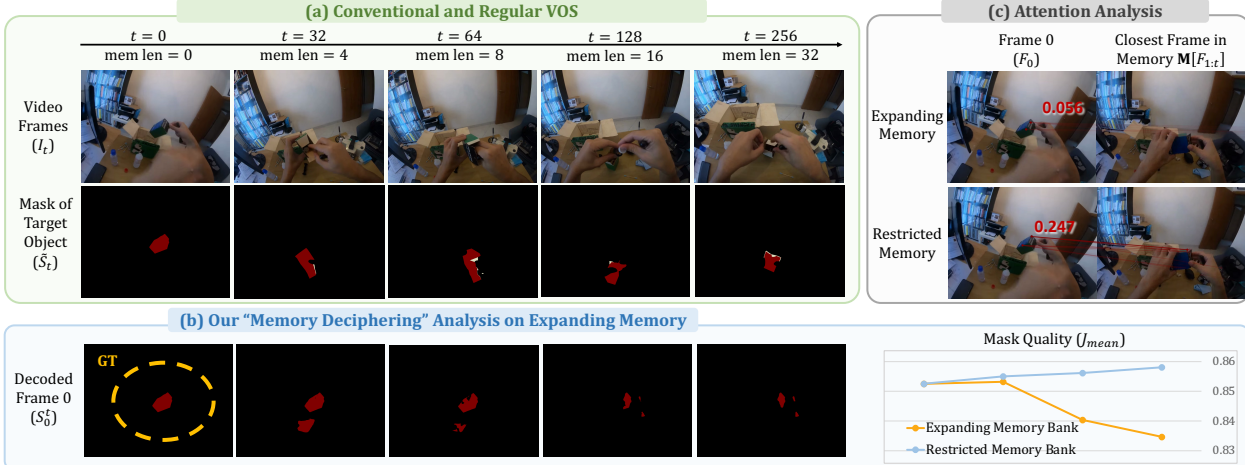
Figure 2. Sketch of Pilot Study. Our *memory deciphering* analysis emulates *decoding the mask on frame 0 from the memory bank features* to quantify the impact of a growing memory bank on VOS modules, where the "desired results" in the figure are the ground truth. For a video shown in Block **(a)**, we visualize its decoding results in Block **(b)**: the masks degrade both quantitatively (yellow curve) and qualitatively, deviating from the desired results. However, selecting a set of concise frames mitigates this issue (blue curve in Block **(b)**). Therefore, we conjecture that the drawback of a growing memory bank lies in confusing the attention of VOS modules. In Block **(c)**, we use red lines to indicate highly weighted associations in attention, with thickness denoting the attention score values. As illustrated, the query F_0 focuses less on its most relevant frame after the memory bank expands, with the attention score dropping from 0.247 to 0.056. (2^{nd} row shows ground-truth masks $\widetilde{S}_t$ as the reference. $\mathcal{J}_{\text{mean}}$ is the average Jaccard between S_0^t and $\widetilde{S}_0$ over all videos.)

3. Pilot Study: Memory Deciphering Analysis

This section devises our pilot experiments on how an expanding memory bank influences the decoding capability of VOS modules. Our design emulates the task of VOS but makes several modifications guided by the principle of controlling variables: the prediction targets and VOS modules are aligned across our pilot experiments, while only the frames in the memory bank vary. Such a comparison enables a clean analysis and reveals the core insight: VOS modules have limited capability to decode a growing memory bank.

Notation and Formulation of VOS. We consider the existing VOS framework as a memory-based encoder-decoder network: the encoder $\mathbf{E}(\cdot)$ is a visual backbone encoding the image I_t at frame t into the feature F_t; and then, the decoder $\mathbf{D}(\cdot)$ converts F_t into the segmentation S_t via reading the features stored in the memory $\mathbf{M}[F_{0:t-1}]$, as below,

$$F_t = \mathbf{E}(I_t), \quad S_t = \mathbf{D}(F_t, \mathbf{M}[F_{0:t-1}]). \quad (1)$$

Here, $\mathbf{M}[F_{0:t-1}]$ generally comes from saving the features at a certain frequency [11, 26, 50], and the VOS decoder is usually special transformers [41], *e.g.*, LSTT in AOT [50]. The final objective of VOS is to minimize the difference between the predicted mask S_t and ground truth $\widetilde{S}_t$.

Design of Our Memory Deciphering Analysis. Our pilot study separates the variables of the VOS module $\mathbf{D}(\cdot)$ and the prediction target $\widetilde{S}_t$ to clearly analyze the influence of the memory bank $\mathbf{M}[F_{0:t-1}]$ under a controlling variable setting. Therefore, we purposefully design our memory deciphering analysis as *decoding the mask of the initial frame*

(frame 0) from the features stored in the memory bank.

More precisely, our pilot study is formulated as,

$$S_0^t = \mathbf{D}'(F_0, \mathbf{M}[F_{1:t}]), \quad (2)$$

where $\mathbf{D}'(\cdot)$ is an additional VOS decoder trained for the objective in Eqn. 2. In practice, we use the original VOS decoder $\mathbf{D}(\cdot)$ to conduct regular VOS as Eqn. 1, and then employ $\mathbf{D}'(\cdot)$ only for deciphering the mask S_t^0 for frame 0, to avoid influencing the original VOS. $\mathbf{M}[F_{1:t}]$ contains the stored features between frames 1 to t. Note that the feature of frame 0 is excluded from the input $\mathbf{M}[F_{1:t}]$ to avoid $\mathbf{D}'$ from trivially relying on single-frame memory.

Before delving into the experiments, we emphasize our reasons for choosing this formulation. (1) *Presence of relevant information*. The procedure in Eqn. 1 resembles propagating the masks from historical frames to the current frame t, indicating that $\mathbf{M}[F_{1:t}]$ contains the information about the mask at frame 0. Therefore, decoding the mask on frame 0 from $\mathbf{M}[F_{1:t}]$ is not a random guess, but should achieve high-quality results. (2) *Identical prediction target*. Our prediction target remains identical for every frame and varying memory size. (3) *Cooperating with regular VOS*. We utilize $\mathbf{D}'(\cdot)$ as a stand-alone VOS decoder so that the original VOS process remains unchanged and our pilot study can utilize the same memory bank.

Implementation. We select the recent VOST [39] dataset to highlight challenging *object state changes*. Its long video duration and complex scenarios push the limits of VOS decoders in deciphering memory. Then we adopt AOT [50] as the VOS encoder-decoder, a popular baseline and the top
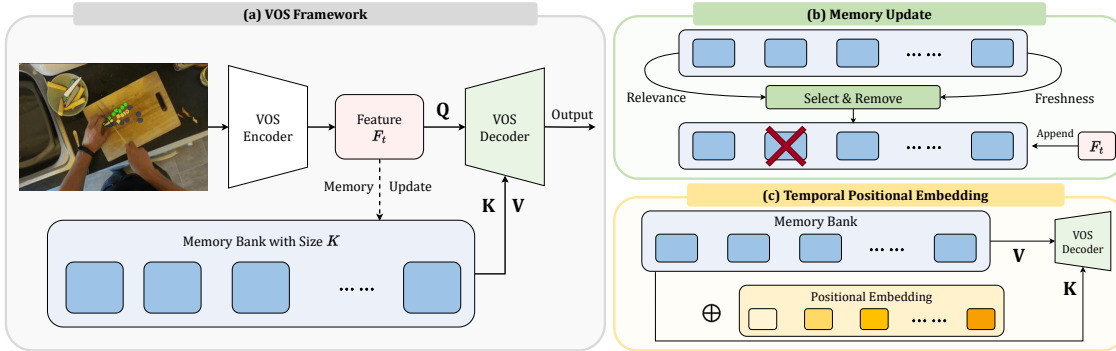
Figure 3. RMem Overview. **(a)** RMem revisits *restricting memory banks* to enhance VOS (Sec. 4.1), motivated by the insight from our pilot study. **(b)** To maintain an informative memory bank, we balance both the relevance and freshness of frames when updating the latest features (Sec. 4.2). **(c)** Benefiting from smaller memory size gaps between training and inference, we introduce previously overlooked temporal positional embedding to encode the orders of frames explicitly (Sec. 4.3), which enhances spatio-temporal reasoning.

method on VOST. Emulating Eqn. 2, we initialize $\mathbf{D}'(\cdot)$ from AOT's pretrained decoder $\mathbf{D}(\cdot)$, and then supervise S_t^0 with a segmentation loss between the ground truth $\widetilde{S}_0$. More implementation details are in Sec. B.

Hypothesis and Expectations. With an expanding memory bank, the information in $\mathbf{M}[F_{1:t}]$ becomes *rigorously richer* at later frames while the prediction target is unchanged. Therefore, we naturally expect the decoded mask S_0^t to illustrate stable or better accuracy at later frames, assuming that the VOS decoder $\mathbf{D}(\cdot)$ is capable of extracting the relevant features from an increasingly large $\mathbf{M}[F_{1:t}]$.

Results and Analysis. Contrasting the expectation above, we observe that masks S_0^t degrade with a growing memory bank, as shown in Fig. 2 (b). To verify that the growing memory bank is indeed the cause of degradation, we empirically bound the memory bank to 8 frames containing the most relevant and latest information, intuitively: first 7 frames and the latest frame in $\mathbf{M}[F_{1:t}]$. According to the blue curve in Fig. 2 (b), *restricting the memory only to store concise features* effectively avoids degradation.

Inspired by addressing the degradation issue, we propose that the *redundant information* is the main negative impact of an expanding memory bank. Otherwise, the degradation should not disappear simply after we select a subset of intuitively relevant frames. More specifically, this closely relates to how VOS methods utilize attention mechanisms to read from memory banks, where the redundant features decrease the attention scores on relevant frames. As direct evidence, we analyze the attention scores for decoding S_0^t in Fig. 2 (c) and observe that the attention scores between F_0 and its most relevant memory feature (first frame in $\mathbf{M}[F_{1:t}]$) have worse concentration on the correct object and become scattered in a longer memory bank. Therefore, we conclude that restricting the memory banks with a concise set of relevant features potentially benefits the decoding of VOS modules via more precise attention.

4. Method of RMem

Motivated by our insight from the pilot study, we propose a straightforward approach highlighting a concise memory bank: restricting the memory with a constant frame number (Sec. 4.1). We then explore the strategies to update the memory bank to constantly digest incoming features and remove obsolete frames (Sec. 4.2). Finally, the restricted memory bank decreases the gap between the memory lengths across the training and inference stages. This enables previously overlooked techniques, and we propose a compelling example of temporal positional embedding (Sec. 4.3). The overview of our method "RMem" ("R" for "Restricted") is in Fig. 3.

4.1. Restricting Memory Banks for VOS

Design. As indicated in our pilot study (Sec. 3), VOS modules have limited capability to process large quantities of features and thus benefit from a concise memory bank with less redundant information. To verify this in actual VOS systems, we develop the simple approach of *restricting the memory bank to a fixed frame number*. In practice, a pre-defined small constant number K is the maximum number of frames a memory bank can store, as shown in Fig. 3. The simplicity of our approach makes it a *plug-and-play* enhancement for the existing VOS framework.

At an arbitrary frame t, we simplify the notation of the memory bank by denoting $\mathbf{M}[F_{0:t-1}]$ as $\mathbf{M}^t$, containing $K_t \leq K$ frames. A natural issue of bounded memory $\mathbf{M}^t$ is that K_t can reach the limit K at sufficiently large t, making the digestion of newly arriving features non-trivial, especially when the quality of information is vital for VOS, according to how we address degradation in the pilot study (Sec. 3). Our baseline adopts an intuitively simple yet effective approach (we explore better strategies in Sec. 4.2): selecting the most reliable frame (frame 0) and temporally most relevant frames (closest frames). Formally, updating the memory bank is as below when $K_t = K$:

$$\mathbf{M}^{t+1} = \text{Concat}(\mathbf{M}_0^t, \ \mathbf{M}_{2:K_t-1}^t, \ F_t), \quad (3)$$

where $\mathbf{M}_{2:K_t-1}^t$ and F_t are the closest frames, and $\mathbf{M}_1^t$ is removed to create an available slot, as shown in Fig. 3 (b).

Discussion. Our restricted memory is a revisit to previous methods [26, 27]. However, we are distinct in emphasizing *accuracy* instead of *efficiency*. In addition, our RMem also simplifies them [9, 26, 27] by treating each frame as a constituent feature map instead of breaking it into smaller regions or pixels [9]; thus, our strategy can directly apply to a wider range of models. Although more sophisticated strategies might further improve our accuracy, a simple approach is already effective (Sec. 5.3).

4.2. Memory Update

Updating the incoming frames to the memory bank provides informative cues for VOS modules to decode. Although our baseline (Eqn. 3) has already cooperated with the bounded memory bank, we investigate better methods for updating.

Challenges of Memory Update. As shown in our pilot study (Sec. 3), improving the conciseness of information heavily influences the decoding efficacy of VOS modules. Therefore, naive heuristics of random selection or keeping the latest frames are unreliable (as in Sec. 5.4, memory update analysis), since they fail to consider the relevance of frames (random) or suffer from drifting of knowledge (latest). To this end, we propose the principles that consider both *relevant* prototypical features and *fresh* incoming information from the latest frames.

Memory Update Inspired by Multi-arm Bandits. Our memory update problem can be stated as *how to select and delete the most obsolete frame k_d from K candidates* to create slots for incoming features. Although not exactly identical, this problem analogizes *multi-arm bandit* [23], which also concerns optimizing the reward by selecting from a fixed number of candidates. Its most inspiring insight for us is balancing the exploitation and exploration with the upper confidence bound (UCB) algorithm [3], whose maximization objective O_k for an option k is as below,

$$O_k = R_k + \sqrt{(2\log T)/t_k}, \qquad (4)$$

where R_k is option k's average reward, T is the total timestamps, and t_k is the number of timestamps selecting k. When applying to our VOS, we re-define R_k as the relevance of a frame for reliable VOS and consider $\sqrt{(2\log T)/t_k}$ as the freshness of memory, intuitively. Then, the deleted frame k_d is chosen according to the smallest $O_{1:K}$. In practice, we define the relevance term R_k using the attention scores between frame $\mathbf{M}_k^t$ and current VOS target F_t, to quantify the contribution of features from the memory. Under the context of transformers, we assume decoding the memory bank is as

$$F_t^D = \mathrm{Attn}(\mathbf{Q} = F_t, \mathbf{K} = \mathbf{M}^t, \mathbf{V} = \mathbf{M}^t), \qquad (5)$$

and assume that $\mathbf{S}^t$ is the scores (after softmax) between F_t and $\mathbf{M}^t$, computed inside the attention. Then, we treat the sum of scores as the relevance of a frame in the memory: $R_k = \mathrm{sum}(S_k^t)$, where S_k^t is the slice of attention scores corresponding to $\mathbf{M}_k^t$. Compared to XMem [9], which also uses attention scores for selection, our design differs in selecting at the frame level instead of the pixel level, which is simpler and already effective (as in Sec. 5.4).

As for the second term in UCB, $\sqrt{(2\log T)/t_j}$, we modify it by defining t_j as the times a frame has stayed in the memory bank and T as the sum of all the frames' staying time. This freshness term penalizes long-staying frames and allows refreshing from the latest information. Finally, O_k combines it with the relevance term R_k via a weight α balancing their numerical scales.

4.3. Memory with Temporal Awareness

Motivation. In addition to accommodating the decoding capability of VOS modules, restricting the memory bank systematically decreases the training-inference discrepancies in memory lengths. Specifically, the VOS algorithms are generally trained on short video clips with a few frames in the memory, while the videos are much longer during inference time. Therefore, the number of frames in the memory bank diverges more significantly without our restriction.

Such temporal alignment between training and inference opens new opportunities for VOS. As a compelling example, we introduce temporal positional embedding (PE) to enhance spatio-temporal reasoning. Specifically, we notice that previous approaches [9, 11, 50] overlook the order of frames in the memory, *i.e.*, the temporal relationship among the frames are not explicitly considered, while spatial PE is widely adopted. Considering the vital role of orders in temporal modeling, which is commonly addressed with temporal PE in video-based tasks, we conjecture that the distinction of memory sizes between training and inference hinders previous methods from employing temporal PE.

Design. The objective of temporal PE is to embed explicit temporal awareness into memory and guide the attention in Eqn. 5. Although restriction on the memory bank alleviates the training-inference shift, the challenges of temporal PE still exist: the optimal memory size K, though much smaller than expanding, can still be larger than the training-time memory size K_{train}; (2) the frames in the memory are varying from 1 to K. To address them, our solution is inspired by how ViT [16] uses learnable PE and interpolation to address different image resolutions. Similarly, we initialize the PE according to K_{train}, denoted as $\widetilde{P}_{0:K_{\text{train}}-1}$, and the query F_t having a dedicated PE P_q. Then, the temporal PE for the memory bank $\mathbf{M}_{0:K_t-1}^t$ is $P_{0:K_t-1}^t$.

$$P_{0:K_t-1}^t = \begin{cases} \widetilde{P}_{0:K_t-1}, & K_t \leq K_{\text{train}} \\ \mathrm{Interp}(\widetilde{P}_{0:K_{\text{train}}-1}, K_t), & K_t > K_{\text{train}} \end{cases} \qquad (6)$$

where "Interp($\cdot$)" interpolates $\widetilde{P}_{0:K_{\text{train}}-1}$ to K_t via nearest neighbor. Finally, temporal PE enhances the original atten-

tion in Eqn. 5 by augmenting the key and values, identical to our conceptual illustration in Fig. 3 (c):

$$F_t^D = \text{Attn}(\mathbf{Q} = F_t + P_q,$$
$$\mathbf{K} = \mathbf{M}_{0:K_t-1}^t + P_{0:K_t-1}^t, \qquad (7)$$
$$\mathbf{V} = \mathbf{M}_{0:K_t-1}^t).$$

The above design contains two critical choices. (1) We use the relative index $\{k = 0, ..., K_t-2\}$ inside the memory instead of the frame index t to avoid the shift between training and inference. (2) Using learnable PE instead of Fourier features fits better to a limited training length, K_{train}.

5. Experiments

5.1. Datasets and Evaluation Metrics

VOST. We primarily utilize the recent VOST [39] dataset that concentrates on challenging object state changes. It curates over 700 videos covering diverse object state changes, *e.g.*, changing appearance, occlusions, crowded objects, and fast motion. In VOST, the evaluation metrics are $\mathcal{J}$ and $\mathcal{J}_{\text{tr}}$, resembling the average Jaccard over *all the frames* and the harder *last 25% frames* corresponding to state changes.

Long Videos Dataset. We use the Long Videos dataset [27] to evaluate long-term understanding, similar to XMem [9]. It contains 3 validation videos with more than 1k frames. $\mathcal{J}$, $\mathcal{F}$ (boundary F measure), and $\mathcal{J}\&\mathcal{F}$ (average of $\mathcal{J}$, $\mathcal{F}$) are considered for evaluation.

LVOS. We also experiment with the recent LVOS [20] dataset and include the results in Sec. C.5.

Regular and Short Video Datasets. YoutubeVOS [45] and DAVIS [35, 36] are two earlier datasets with *shorter* duration and *easier* scenarios compared with VOST. In this paper, we use them as the pretraining datasets for VOST and the Long Videos dataset following standard practice [9, 39], and conduct analysis in addition to the challenging datasets.

5.2. Baselines and Implementation Details

Our proposed RMem is a simple and plug-and-play enhancement for the VOS framework. Without loss of generality, we select AOT [50] and DeAOT [48] as the main baseline because of its top performance on VOST (as in Table 1) and simplicity. It adopts ResNet-50 [19] as its encoder and a specially designed "long short term-transformer" (LSTT) as its decoder. For the memory bank, the original AOT digests the latest frame and expands the memory continuously, while RMem restricts its size to 8 frames. We also employ RMem on other VOS methods in addition to AOT. More details on models and implementation in Sec. B.

5.3. State-of-the-art Comparisons

VOST. In Table 1, we compare RMem with previous methods on VOST. Our approach establishes new state of the art on this challenging benchmark with a significant improvement. Notably, our simple strategy increases the VOS

	$\mathcal{J}_{\text{tr}}$	$\mathcal{J}$
OSMN Match [47]	7.0	8.7
OSMN Tune [47]	17.6	23.0
CRW [22]	13.9	23.7
CFBI [49]	32.0	45.0
CFBI+ [51]	32.6	46.0
XMem [9]	33.8	44.1
HODOR Img [1]	13.9	24.2
HODOR Vid [1]	25.4	37.1
AOT [50]	36.4	48.7
AOT$^\Psi$	37.0	49.2
AOT$^\Psi$ + RMem (Ours)	39.8	50.5
DeAOT$^\Psi$	37.6	50.1
DeAOT$^\Psi$ + RMem (Ours)	**40.4**	**51.8**

Table 1. Comparisons with previous methods on VOST [39]. Our RMem shows advantages on both overall quality ($\mathcal{J}$) and addressing object state changes ($\mathcal{J}_{\text{tr}}$). (If not specified, the results are from VOST's implementation, Ψ denotes our implementation.)

	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
CFBI [49]	53.5	50.9	56.1
CFBI+ [51]	50.9	47.9	53.8
STM [34]	80.6	79.9	81.3
MiVOS [10]	81.1	80.2	82.0
AFB-URR [27]	83.7	82.9	84.5
STCN [11]	87.3	85.4	89.2
XMem [9]	89.8	88.0	91.6
AOT [50]	84.3	83.2	85.4
AOT$^\Psi$	86.7	85.5	87.9
AOT$^\Psi$ + RMem (Ours)	90.3	88.5	92.1
DeAOT$^\Psi$	89.4	87.4	91.4
DeAOT$^\Psi$ + RMem (Ours)	**91.5**	**89.8**	**93.3**

Table 2. Comparison with previous methods on Long Videos dataset [27]. For both baselines of AOT and DeAOT, our RMem shows significant improvement. (Without mention, the results are from XMem [9], Ψ denotes our implementation.)

quality for the whole video ($\mathcal{J}$) and maintains better robustness for the state-changing frames ($\mathcal{J}_{\text{tr}}$). This is especially clear when compared to AOT [50]: the improvement is over $\sim$3% with our plug-and-play modifications.

Long Videos Dataset. As our RMem limits memory capacity, a natural suspicion is that our memory bank performs worse in storing information and struggles with long-term modeling. However, our comparison in Table 2 shows the opposite. On the Long Videos dataset, our RMem not only improves upon the baseline AOT and DeAOT models but also outperforms the state of the art XMem [9] model, which utilizes specially designed hierarchical memory banks and memory manipulation operators. Therefore, this further supports our insight on keeping a concise memory bank to accommodate the limited capability of VOS modules to address expanding memory banks.

5.4. Ablation Studies

Effect of RMem Components. We analyze each RMem component respect to AOT and DeAOT baselines, as in Table 3. **(1)** *Restricting memory banks.* The most important insight from our pilot study (Sec. 3) is to maintain a concise memory bank with relevant information, which motivates our revisit of restricting memory banks (Sec. 4.1). Accord-

ID	RM	TPE	MU	AOT $\mathcal{J}_{tr}$	AOT $\mathcal{J}$	DeAOT $\mathcal{J}_{tr}$	DeAOT $\mathcal{J}$
1	Baseline			37.0	49.2	37.6	50.9
2	✓			38.7	50.3	38.8	51.0
3	✓	✓		39.7	50.3	40.0	51.7
4	✓		✓	39.4	50.3	39.0	51.4
5	✓	✓	✓	**39.8**	**50.5**	**40.4**	**51.8**

Table 3. Ablation studies of RMem on VOST. Starting from the AOT and DeAOT baselines, all of the components improve the performance, especially the harder object state-changing frames ($\mathcal{J}_{tr}$). **RM**: restricting memory banks. **TPE**: temporal positional embedding. **MU**: memory update with the UCB algorithm.

Method	Variant	$\mathcal{J}_{tr}$	$\mathcal{J}$
Remove	0^{th}	35.9	48.9
	1^{st}	38.7	50.3
	Middle	38.3	50.2
	Latest	35.7	48.5
	Random	38.0	50.0
UCB	Relev	39.1	50.1
	Relev + Fresh	**39.4**	**50.3**

Table 4. Ablation study of different memory updating strategies on VOST. We analyze deleting a frame in the memory based on heuristics ("Remove") or guided by the relevance and freshness of the UCB algorithm ("UCB"). Our final memory updating strategy using both relevance and freshness achieves the best performance.

ing to Table 3 (row 1 and 2), a bounded memory bank leads to significant enhancement in the long and complex VOST videos. **(2)** *Temporal positional embedding.* In Table 3, we illustrate that adding positional embedding (Sec. 4.3) greatly benefits the spatio-temporal modeling, especially the harder $\mathcal{J}_{tr}$ for state changes. **(3)** *Memory update.* We refresh the memory banks by balancing the relevance and freshness of frames (Sec. 4.2), inspired by the UCB algorithm [3]. In rows 4 and rows 5 of Table 3, such a strategy effectively boosts the overall performance.

Analysis on Frame Numbers of Memory Banks. We verify a direct implication of our insight: an expanding memory bank elevates the difficulty of VOS modules to decode information. Specifically, we observe the VOS accuracy under various sizes of memory banks. To avoid the influence of hyper-parameter tuning, we utilize the baseline memory update strategy in Sec. 4.1. As in Fig. 4, the performance first improves from richer information. Then both $\mathcal{J}$ and $\mathcal{J}_{tr}$ decrease when the length of memory exceeds the capability of learned AOT modules, until they become similar to unrestricted memory. Consequently, these results directly support our insight of restricting memory banks.

Memory Update Analysis. Maintaining an informative memory bank is critical for the VOS accuracy, and we propose a UCB-inspired algorithm in Sec. 4.2. Table 4 analyzes the key intuition and design choices with AOT. **(1)** The initial frame is critical in keeping the provided ground-truth information: removing the 0-th frame leads to an accuracy drop, and is more profound when scenarios are complex (VOST). **(2)** Guaranteeing the freshness of information is
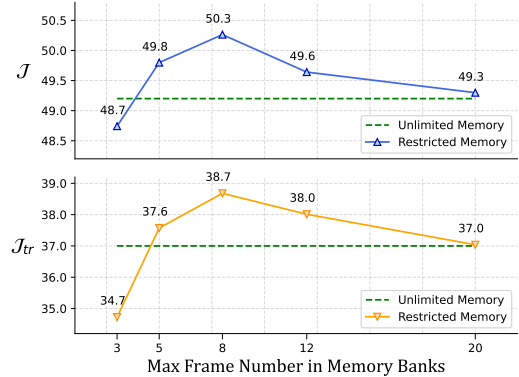

Figure 4. Impact of memory bank size on VOS, tested on VOST. With more frames in the restricted memory, the accuracy first increases and then decreases until it approximates unrestricted memory. This supports the limited deciphering capability of VOS modules and our insight into restricting memory banks.

Method	$\mathcal{J}_{tr}$	$\mathcal{J}$
AOT	37.0	49.2
AOT + RM	38.7	50.3
AOT + SinCos PE	37.2	48.3
AOT + Learnable PE	36.7	49.4
AOT + RM + SinCos PE	37.9	48.9
AOT + RM + Learnable PE	**39.7**	**50.3**

Table 5. Comparison of temporal PE strategies on VOST. Based on restricted memory ("RM"), our learnable temporal PE ("Learnable") is better than using high-frequency Fourier features ("Sin-Cos"). Notably, restricting memory is essential for PE.

critical, where removing the latest frame leads to the worst accuracy. **(3)** Randomly removing frames performs surprisingly well but is still worse than our baseline (removing the 1^{st} frame, in Sec. 4.1). **(4)** Using attention scores to reflect the relevance better removes redundant features ("Relev"), and it is further enhanced with the freshness term, where freshness is especially effective to avoid frames from staying long time in the memory bank, supported by the Long Video dataset. Finally, the best strategy is our UCB-inspired algorithm combining relevance and freshness.

Temporal Positional Embedding Strategies. We introduce using learnable temporal PE to address the varied frames in the memory banks of VOS in Sec. 4.3. In Table 5, we analyze another PE strategy of encoding the index into high-frequency features with SinCos functions and find it performs worse. This is because SinCos is commonly used in scenarios of a large number or continuous space of coordinates (*e.g.*, NeRF [31]), while learnable embeddings can better handle a small number of slots (*e.g.*, ViT [16]), as in the limited memory length during the VOS training. Furthermore, we highlight that temporal PE requires restricted memory to function well because of better training-inference temporal alignment in memory lengths. This supports our intuition in Sec. 4.3 and suggests the emerging opportunities from restricting memory banks.
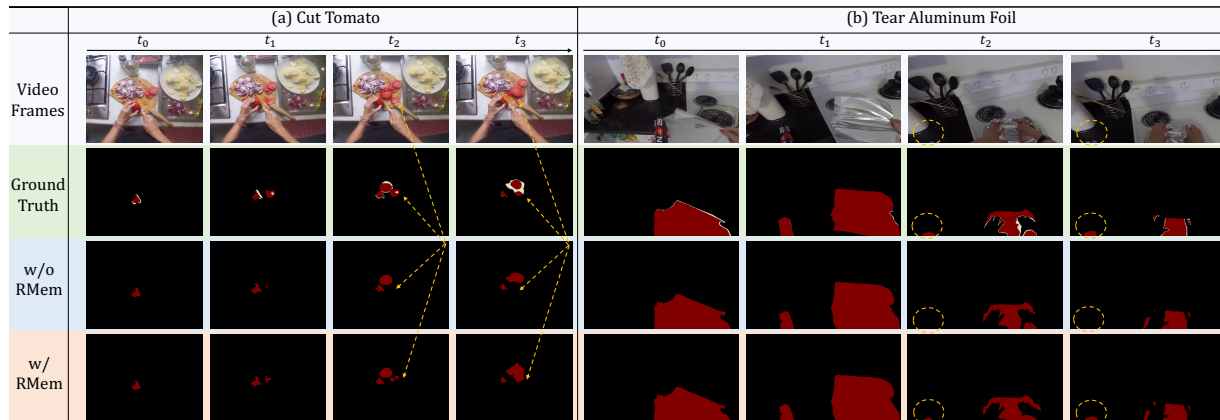
Figure 5. (Best viewed zoom-in with color.) Qualitative VOS results for object state changes on VOST [39]. We provide two examples showing the challenges of object state changes, including slicing, occlusions, distraction from similar objects (other tomatoes), and shape changes. For both scenarios, using RMem shows advantages in robustly maintaining the masks of the target objects, as highlighted. (White pixels are annotated by VOST denoting "ignored" regions for evaluation, which are hard and ambiguous even for human annotators.)

Method	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}_{tr}$	$\mathcal{J}$	Max Mem $\downarrow$	FPS
AOT	85.2	82.5	87.9	4.46G	13.67
AOT + RMem (Ours)	85.2	82.4	88.0	2.34G	15.57
DeAOT	85.2	82.3	88.1	2.24G	25.11
DeAOT + RMem (Ours)	**85.3**	**82.4**	**88.2**	**1.53G**	**27.42**

Table 6. RMem maintains the accuracy on DAVIS2017 while being more efficient, indicating that RMem can be generally applied, not limited to challenging scenarios. This also aligns with the prior works and suggests that not having demanding datasets was potentially why the *accuracy* benefits of memory restriction were not clearly revealed previously.

Analysis on Regular and Short Video Benchmarks. We highlight the improvement on long and complex VOS datasets, but we also supplement our analysis on the regular and short video dataset DAVIS2017. As in Table 6, our RMem has relatively the same performance but effectively improves the efficiency. Compared with our improvement on VOST and the Long Video dataset, we conjecture that the learned VOS modules (AOT and DeAOT) are already capable of handling shorter video duration and less complicated scenarios, even without our concise memory banks. Additionally, this potentially explains that previous studies exploring restricting memory banks [25, 27] have not explicitly discovered its benefits, *probably due to not considering longer and more challenging datasets like VOST.*

5.5. Qualitative Results

We visualize on two representative videos from VOST [39] that require robust spatio-temporal reasoning in Fig. 5. Video (a) is the kitchen behavior of cutting a tomato into slices, and it illustrates the challenges of splitting objects, occlusions from hands, and visual distraction from other tomatoes. Without our RMem, the baseline AOT model fails to maintain the masks for the separated tomato slice, while using RMem correctly remembers this slice at the later stage of the video (columns 3 and 4). Such regions are highlighted with the yellow arrows. The other video (b)

illustrates another difficulty of object shape transformation and splitting between the box and the aluminum. Although the baseline model without RMem can correctly segment the box at the beginning of splitting (column 2), it gradually loses track of the box and can only concentrate on the dominant object. However, our model enhanced with RMem robustly segments the small regions of the box, indicating that its attention association with relevant historical frames is still stable because of our restricted memory. Therefore, we conclude that the quantitative results reveal the difficulties of object state changes and support the effectiveness of our approach.

6. Conclusion

This paper reveals the drawbacks of expanding memory banks, a conventional design in VOS. Our insight stems from a novel "memory deciphering" analysis, which suggests that the redundant information in growing memory banks confuses the attention of VOS modules and elevates the difficulty of feature decoding. Then, we propose the simple enhancement for VOS named RMem. At its core is restricting the size of memory banks, accompanied by UCB-inspired memory update strategies and temporal positional embedding to enhance spatio-temporal reasoning. Extensive evaluation on the recent challenging datasets, including VOST and the Long Videos dataset, supports our insight and effectiveness of RMem.

Limitations and Future Work. Our paper prioritizes the analysis of memory banks and illustrates our insight with a straightforward approach. Therefore, interesting future work is to combine the intuition from more sophisticated methods, such as XMem [9]. Furthermore, our exploration mainly adapts memory banks to cooperate with the capability of VOS modules, while how to improve the decoding ability of VOS modules for a huge memory bank is the alternative direction and interesting future work.

References

[1] Ali Athar, Jonathon Luiten, Alexander Hermans, Deva Ramanan, and Bastian Leibe. HODOR: High-level object descriptors for object re-segmentation in video learned from static images. In *CVPR*, 2022. 2, 6, 14

[2] Ali Athar, Alexander Hermans, Jonathon Luiten, Deva Ramanan, and Bastian Leibe. TarVis: A unified approach for target-based video segmentation. In *CVPR*, 2023. 2

[3] Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *JMLR*, 3(Nov):397–422, 2002. 2, 5, 7, 12

[4] 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 *ECCV*, 2020. 2

[5] Thomas Brox and Jitendra Malik. Object segmentation by long term analysis of point trajectories. In *ECCV*, 2010. 2

[6] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *CVPR*, 2017. 2

[7] Yuhua Chen, Jordi Pont-Tuset, Alberto Montes, and Luc Van Gool. Blazingly fast video object segmentation with pixel-wise metric learning. In *CVPR*, 2018. 2

[8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022. 2

[9] Ho Kei Cheng and Alexander G Schwing. XMem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *ECCV*, 2022. 1, 2, 5, 6, 8, 12, 13, 14, 15

[10] 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*, 2021. 6

[11] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. In *NeurIPS*, 2021. 1, 2, 3, 5, 6

[12] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. In *CVPR*, 2024. 2

[13] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The EPIC-KITCHENS dataset. In *ECCV*, 2018. 2

[14] Ahmad Darkhalil, Dandan Shan, Bin Zhu, Jian Ma, Amlan Kar, Richard Higgins, Sanja Fidler, David Fouhey, and Dima Damen. EPIC-KITCHENS VISOR benchmark: Video segmentations and object relations. In *NeurIPS*, 2022. 2

[15] 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. In *ICCV*, 2023. 2

[16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 5, 7

[17] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W Taylor. SSTVos: Sparse spatiotemporal transformers for video object segmentation. In *CVPR*, 2021. 2

[18] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4D: Around the world in 3,000 hours of egocentric video. In *CVPR*, 2022. 2

[19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6, 11

[20] Lingyi Hong, Wenchao Chen, Zhongying Liu, Wei Zhang, Pinxue Guo, Zhaoyu Chen, and Wenqiang Zhang. LVOS: A benchmark for long-term video object segmentation. In *ICCV*, 2023. 2, 6, 14

[21] Yuan-Ting Hu, Jia-Bin Huang, and Alexander G Schwing. Videomatch: Matching based video object segmentation. In *ECCV*, 2018. 2

[22] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. In *NeurIPS*, 2020. 6

[23] Michael N Katehakis and Arthur F Veinott Jr. The multi-armed bandit problem: decomposition and computation. *Mathematics of Operations Research*, 12(2):262–268, 1987. 5

[24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2014. 12

[25] Fuxin Li, Taeyoung Kim, Ahmad Humayun, David Tsai, and James M Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*, 2013. 2, 8

[26] Yu Li, Zhuoran Shen, and Ying Shan. Fast video object segmentation using the global context module. In *ECCV*, 2020. 2, 3, 5

[27] Yongqing Liang, Xin Li, Navid Jafari, and Jim Chen. Video object segmentation with adaptive feature bank and uncertain-region refinement. In *NeurIPS*, 2020. 1, 2, 5, 6, 8, 11, 12, 13

[28] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 11

[29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 12

[30] Kevis-Kokitsi Maninis, Sergi Caelles, Yuhua Chen, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. Video object segmentation without temporal information. *TPAMI*, 41(6):1515–1530, 2018. 2

[31] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 7

[32] Sebastian Nowozin. Optimal decisions from probabilistic models: the intersection-over-union case. In *CVPR*, 2014. 11

[33] Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, and Seon Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, 2018. 2

[34] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *ICCV*, 2019. 6

[35] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 2, 6

[36] 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. 2, 6, 11, 12, 14

[37] Andreas Robinson, Felix Jaremo Lawin, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. Learning fast and robust target models for video object segmentation. In *CVPR*, 2020. 2

[38] Hongje Seong, Junhyuk Hyun, and Euntai Kim. Kernelized memory network for video object segmentation. In *ECCV*, 2020. 2

[39] Pavel Tokmakov, Jie Li, and Adrien Gaidon. Breaking the "object" in video object segmentation. In *CVPR*, 2023. 1, 2, 3, 6, 8, 11, 12

[40] David Tsai, Matthew Flagg, Atsushi Nakazawa, and James M Rehg. Motion coherent tracking using multi-label mrf optimization. *IJCV*, 100(Dec):190–202, 2012. 2

[41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2, 3, 11

[42] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for video object segmentation. In *BMVC*, 2017. 2

[43] 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 *CVPR*, 2019. 2

[44] Haozhe Xie, Hongxun Yao, Shangchen Zhou, Shengping Zhang, and Wenxiu Sun. Efficient regional memory network for video object segmentation. In *CVPR*, 2021. 2

[45] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-VOS: Sequence-to-sequence video object segmentation. In *ECCV*, 2018. 2, 6, 11, 12, 14

[46] Zihui Xue, Kumar Ashutosh, and Kristen Grauman. Learning object state changes in videos: An open-world perspective. In *CVPR*, 2024. 1

[47] Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, and Aggelos K Katsaggelos. Efficient video object segmentation via network modulation. In *CVPR*, 2018. 2, 6

[48] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. In *NeurIPS*, 2022. 2, 6, 11, 14

[49] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *ECCV*, 2020. 2, 6, 12

[50] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. In *NeurIPS*, 2021. 1, 2, 3, 5, 6, 11, 14

[51] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by multi-scale foreground-background integration. *TPAMI*, 44(9):4701–4712, 2021. 2, 6

[52] Jiangwei Yu, Xiang Li, Xinran Zhao, Hongming Zhang, and Yu-Xiong Wang. Video state-changing object segmentation. In *ICCV*, 2023. 1, 2

Appendix

Our appendix cover additional analysis, implementation details, and discussion as below:

(A) **Demo Video.** We provide a demo video at https://youtu.be/mFjGSPXmXdA showing multiple challenging VOS examples (Sec. A).

(B) **Implementation details.** We explain the detailed model architectures and the procedures for training and inference (Sec. B).

(C) **Additional ablation studies.** This section provides more analysis and experimental results (Sec. C).

(D) **Addition discussion on limitations and future work.** We offer a more detailed discussion of the limitations and potential future directions (Sec. D).

A. Demo Video

In https://youtu.be/mFjGSPXmXdA, we provide four qualitative comparison examples between the baseline models (AOT [50] and DeAOT [48]) and our RMem, with the object state changes from both VOST [39] and the Long Videos dataset [27]. Notably, these examples illustrate four challenging scenarios: (1) **Object ambiguity**: objects have similar appearances; (2) **Slicing**: an object is cut into multiple slices; (3) **Appearance changes**: an object has changed its shape and appearances, leading to incorrect VOS masks. (4) **Sudden shape changes**: the viewpoint changes quickly and causes variation in shapes of the target object. The four examples demonstrate that RMem effectively improves the spatio-temporal reasoning of VOS.

B. Implementation Details

We describe the outline of implementation of AOT [50] and DeAOT [48] baselines in Sec. 5.2. This section provides in-depth details of the implementation and the configuration of RMem.

B.1. Model Architecture

AOT and DeAOT share the common architecture of the memory-based VOS framework. As conceptualized in Eqn. 1, we disassemble the VOS framework into the modules of an encoder $\mathbf{E}(\cdot)$ encoding images into feature maps, a decoder $\mathbf{D}(\cdot)$ extracting information from the memory bank, and a segmentation head translating the output from decoder into masks. Please note that we have additionally decoupled the segmentation head from the decoder for clarity, compared with Eqn. 1.

Encoder. Identical to VOST [39], we adopt ResNet-50 [19] as the encoder, which achieves competitive performance while efficient enough to operate on Long Videos. The multiple stages in the ResNet encoder produce 3 levels of feature maps $\{F^4, F^8, F^{16}\}$ with 1/4, 1/8, and 1/16

the resolution of the original input image, respectively. Following the practice of AOT and DeAOT, the deepest feature map F^{16} is the input to the decoder for memory reading, and $\{F^4, F^8\}$ are provided to the segmentation head as input for predicting high-quality masks.

Decoder. AOT and DeAOT utilize a specially-designed transformer [41] to conduct associative memory reading, named "Long Short-term Transformer" (LSTT). LSTT comprises three consecutive transformer layers to enhance features in the current frame with the memory bank. Adopting the same notations as Eqn. 1, we conceptually illustrate this process as Eqn. H:

$$
\begin{aligned}
F_t^{(l+1)} = \mathrm{Attn}(\mathbf{Q} &= F_t^{(l)}, \\
\mathbf{K} &= \mathbf{M}^{(l)}[F_{0:t-1}], \\
\mathbf{V} &= \mathbf{M}^{(l)}[F_{0:t-1}]),
\end{aligned}
\tag{H}
$$

where the superscript (l) denotes the layer index of LSTT, ranging from 0 to 2. After the above process, We keep the implementation details identical to the original AOT and DeAOT. Please refer to them for more detailed configuration. Finally, the output feature $F_t^{(3)}$ replaces the feature map F^{16} from the encoder not enhanced with spatio-temporal information.

Segmentation Head. To maintain high-resolution segmentation masks, the segmentation process involves a feature pyramid network (FPN) [28]. It accepts $F_t^{(4)}$ as the input feature, uses $\{F^8, F^{16}\}$ as shortcut inputs, and up-samples them via the combination of a convolutional layer and a bi-linear up-sampling layer.

Temporal Positional Embedding. We introduce temporal positional embedding (TPE) in Sec. 4.3 to enhance the spatio-temporal reasoning ability of models. In practice, we initialize end-to-end learnable embeddings with the same number to the memory length during the training time (*e.g.*, 4 in VOST) and the same dimension to the feature F_t, marking the PE of each place in the memory bank. For simplicity, the three LSTT layers in Eqn. H share the same set of TPE.

B.2. Training

Loss Functions. Our training procedure utilizes the same loss functions as AOT and DeAOT: the combination of bootstrapped cross-entropy loss and soft Jaccard loss [32]. Both loss terms are averaged 1:1 as the final loss value.

VOST. The training on VOST [39] follows the original practice of VOST's authors, where the models are fine-tuned on VOST with pretrained weights from DAVIS2017 [36] and Youtube2019 [45]. As VOST highlights spatio-temporal modeling, we follow the authors' implementation of AOT by using a long sequence length of 15 frames during training and this accordingly enables

4 frames in the memory bank. It leverages exponential moving averages (EMA) for parameter updates to stabilize the training process. The whole training process uses AdamW [24, 29] optimizer, and lasts 20,000 steps with a batch size of 8, on 4×A40 GPUs. The initial learning rate is 2×10^{-4} and it gradually decays to 2×10^{-5} according to a polynomial pattern [49]. To avoid overfitting, we set the learning rate of the encoder as 0.1 of the other components. The weight decay is 0.07, which is also identical to AOT and DeAOT.

Long Videos Dataset. Following the standard practice [9, 27], we first train the AOT and DeAOT models on the DAVIS2017 [36] and YoutubeVOS2019 dataset [45], then conduct inference on the Long Videos dataset [27]. However, to support the training of positional embedding, we extend the length of training samples from the original 5 frames to 9 frames, to support 4 frames in the memory banks during the training time. Please note that we also re-train the baselines under the same setup to ensure a fair comparison. The training procedure leverages the similar optimization setting as described above for the VOST dataset, including the AdamW [24, 29] optimizer, weight decay of 0.07, polynomial learning rate decay [49] from 2×10^{-4} to 2×10^{-5}, 0.1 scaling of the encoder learning rate, and EMA parameter updates. The only difference from VOST is training 100,000 steps with a batch size of 16, following the implementation of the original AOT and DeAOT on DAVIS2017 and YoutubeVOS2019 datasets.

B.3. Inference.

VOST. Instead of appending features into memory at *a fixed frequency of 5 frames*, the authors of VOST developed a different strategy than on DAVIS2017 and YoutubeVOS2019 to address the CUDA memory issue caused by higher resolution and longer video duration: the memory bank is bounded by 30 frames and the frequency of updating memory banks is accordingly $L/30$, where L is the length of the video. For our RMem, we follow the frequency of memory updates set by VOST, but bounds the size of memory banks to 9 frames, which is significantly smaller than the original cap of 30 frames. Therefore, our RMem needs to update the memory banks by removing the obsolete frames, and we describe the details of memory update in Sec. B.4 below.

Long Videos Dataset. When comparing to the other approaches on the Long Videos dataset (Table 2), we primarily rely on the VOS performance evaluated by XMem [9]. However, we re-implement the baselines of AOT and DeAOT for a fair comparison with RMem, since XMem has not released the code for evaluating both methods. Notably, *our re-implementation achieves better performance* compared to XMem's reported numbers. In practice, we de-

termine the frequency of updating memory banks by $L/30$ to avoid CUDA memory issues, which is similar to the inference procedure on VOST. Our RMem shares the same inference setting as baseline, only restricting the memory bank size to 8 frames. Then, the memory update strategy is identical to VOST, as described in Sec. B.4.

B.4. Memory Update

As is described in Sec. 4.2, our RMem balances the relevance and freshness of frames in the memory bank using our algorithm inspired by UCB [3].

Relevance. As mentioned in Sec. 4.2, we use the attention scores from the transformers in the decoder Eqn. 5 to reflect the relevance of a memory frame R_k. Since the LSTT decoder in AOT and DeAOT has three transformer layers, we intuitively select the attention scores from the 0-th transformer because it is closest to the original image embeddings F_t and memory features $\mathbf{M}^t$ (ablation in Sec. C.3). To stabilize the relevance term and avoid fluctuations, we further apply the moving average technique to the relevance term. Suppose R_k' denotes the relevance values of a memory frame k derived from the latest timestamp, the consequent relevance term R_k is updated via:

$$R_k \leftarrow (1 - \lambda)R_k' + \lambda R_k, \qquad \text{(I)}$$

where we set $\lambda = 0.8$ for both VOST and the Long Videos dataset. As we have noticed, using moving average for stabilization is a common technique for VOS on long videos, such as in AFB-URR [27].

Freshness. To balance the numerical scales of the relevance and freshness terms, we slightly modify Eqn. 4 as below,

$$O_j = R_j + \alpha \sqrt{\frac{\log T}{t_j + B}}, \qquad \text{(J)}$$

where B smooths the numerical ranges of the freshness term, and α controls the individual contribution of relevance and freshness. In practice, we set $B = 8$ and $\alpha = 1.5$ for both VOST and the Long Videos dataset. Detailed ablation studies on the values of α are illustrated in Sec. C.2.

C. Supplemental Ablation Studies

C.1. Memory Update on the Long Videos Dataset

We analyze the memory update strategies on the Long Videos dataset [27] using our AOT baseline in Table A, in addition to the analysis on VOST [39] (Table 4). **(1)** Notably, we observe consistent improvement from our UCB-inspired memory update strategy combining both relevance and freshness of frames in the memory. **(2)** Similar to the results on VOST, our baseline of removing the 1-st frame in the memory has competitive performance but is inferior to our final UCB-inspired strategy. **(3)** The analysis in Table A also reveals several intriguing differences between the Long

Videos dataset and VOST. Specifically, VOST highly relies on the relevance of frames and the reliable information from the 0-th frames because of its complexity in scenarios, while the Long Videos dataset highlights the utility of freshness of frames as a consequence of extremely long video duration.

Method	Variants	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
Remove	0^{th}	88.1	86.3	89.9
	1^{st}	88.3	86.6	90.1
	Middle	86.6	85.5	87.9
	Latest	85.4	84.1	86.7
	Random	87.7	86.6	88.9
UCB	Relev	86.9	85.4	88.3
	Relev + Fresh	**89.5**	**87.8**	**91.2**

Table A. Ablation study of different memory updating strategies on the Long Videos dataset, in addition to VOST (Table 4). We analyze deleting a frame in the memory based on heuristics ("Remove") or guided by the relevance and freshness of the UCB algorithm ("UCB"). Our final memory updating strategy using both relevance and freshness achieves the best performance.

C.2. Balancing Relevance and Freshness

As mentioned in Sec. 4.2 and Sec. B.4, we balance relevance and freshness when updating the memory banks via Eqn. J. Fig. A analyzes the performance under different α values on both VOST and the Long Videos dataset. Specifically, a larger α denotes relying more on the freshness term. A proper α is essential for the UCB-inspired algorithm to improve memory update for both VOST and the Long Videos dataset, and we empirically select $\alpha = 1.5$ because it generalizes better to both of the datasets. Interestingly, Fig. A also reveals the difference between VOST and the Long Videos dataset: VOST has more complex scenarios and highlights the utility of relevance, while the long video dataset relies more on freshness due to its extremely long video duration. Nonetheless, our final $\alpha = 1.5$ achieves proper balance for both domains.

C.3. Relevance Calculation

Our relevance term for memory update uses attention scores to reflect the importance of a frame, similar to previous works [9, 27]. However, LSTT has three transformer layers and enables two intuitive strategies of relevance calculation: (1) directly using the 0-th layer; and (2) computing the average attention scores of all the transformer layers. Table B compares these two strategies on VOST and the Long Videos dataset. We observe that using the 0-th layer for relevance calculation has an advantage in most of the scenarios. We conjecture that the 0-th transformer has the largest fidelity to the features of images and memory banks. Therefore, our RMem empirically selects the 0-th transformer for relevance, as described in Sec. B.4.
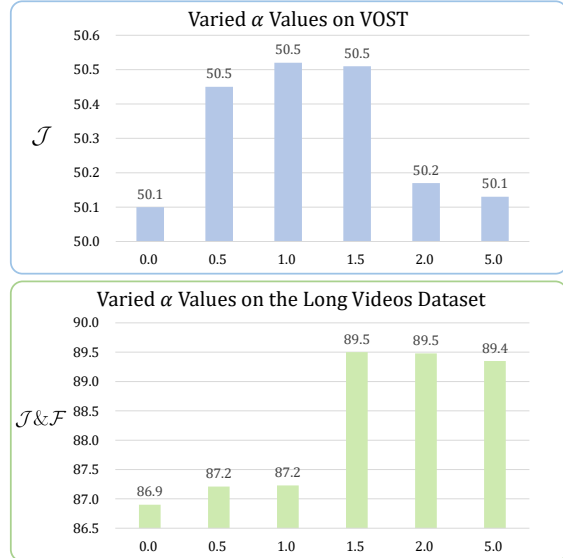

Figure A. Analysis on relevance and freshness for memory update, on VOST and the Long Videos dataset. The performance varies with different α values (from Eqn. J), and it illustrates the importance of the trade-off between relevance and freshness.

Methods	VOST		Long Video		
	$\mathcal{J}_{tr}$	$\mathcal{J}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
AOT + RMem (0^{th})	**39.8**	**50.5**	**90.3**	**88.5**	**92.1**
AOT + RMem (Mean)	39.6	50.3	89.8	88.2	91.5
DeAOT + RMem (0^{th})	40.4	51.8	**91.5**	**89.8**	**93.3**
DeAOT + RMem (Mean)	**40.6**	**52.0**	90.3	88.7	92.0

Table B. Analysis on the relevance calculation. "0-th" and "Mean" denote using the attention scores from the 0-th transformer layer or the average attention scores from all of the three layers.

C.4. Analysis on Training-Inference alignment

As discussed in Sec. 4.3, the purpose of temporal positional embedding is to align the gap between training and inference, as VOS models are trained on short videos but inferencing on unlimited videos. However, it is also valuable to explore whether it is another approach to address this training-inference gap. We compared our Restricted Memory (RM) with 2 approaches: **(1) Longer Memory (LM):** train the model with longer video clips so that the model can fit better on a larger memory bank. **(2) More Steps (MS):** train the model with more steps. As is shown in Table C, LM certainly is effective in mitigating the training-inference gap, but it is still worse than our RM. MS exhibits overfitting with too many training steps, thus not capable of addressing this issue. However, MS can still gain improvement through our RM, proving our method's effectiveness from another perspective.

C.5. Analysis on LVOS

Since the Long Videos dataset only features 3 testing videos, which is not able to fully demonstrate the effective-

Model	Train_Mem_Len	Step	URM		RM	
			$\mathcal{J}_{tr}$	$\mathcal{J}$	$\mathcal{J}_{tr}$	$\mathcal{J}$
AOT	4	20k	37.0	49.2	38.6	50.2
AOT-LM	6	20k	38.2	49.9	39.8	50.1
AOT-MS	4	40k	36.6	48.6	37.8	48.0

Table C. Analysis of 2 approaches to address training-inference gap. "URM" for unrstricted memory and "RM" for restricted memory. Our "RM" is still the best way to align training and inference.

ness of our method, we further report our model's performance on LVOS dataset [20], which contains 50 long videos in the validation set.

Methods	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
AOT	63.6	57.6	69.5
AOT + TPE	64.5	58.9	70.0
AOT + RMem	**66.1**	**60.5**	**71.7**

Table D. Results on the validation set of LVOS dataset.

As is shown in Table D, our RMem still holds the highest performance compared to the AOT baseline. Besides, our TPE (temporal positional embedding) exhibits considerable improvements, which proves that TPE is effective in aligning the training-inference gap, given that the average duration in LVOS is much longer than other video datasets.

C.6. Analysis on YoutubeVOS2019

Our study concentrates on improving the VOS accuracy for long and/or complex VOS scenarios. Meanwhile, we also supplement with analysis on shorter, simpler benchmarks. As indicated in Table 6, our RMem demonstrates comparable performance to baselines without RMem on DAVIS2017 [36], with a notable increase in efficiency. This result underlines the adaptability of our approach across different regimes.

Further analysis is conducted in the section using the YoutubeVOS2019 [45] benchmark, with shorter video duration and easier scenarios. In Table E, we evaluate two settings: (1) the influence of only restricting the memory bank sizes; and (2) the effect of the full RMem with temporal positional embedding. Table E (rows 1 and 2) shows that: by limiting the memory banks with the *original checkpoint* provided by DeAOT's authors, we maintain the same VOS quality. This finding suggests that *constraining the memory banks is a regime-independent strategy*.

A key aspect of our RMem is temporal positional embedding (TPE), which necessitates end-to-end model training on extended sequences. As in Sec. B.2, we *increase the training sequence length from 5 frames to 9 frames* without tuning the hyper-parameters, ensuring a 4-frame memory bank during the training stage. However, this introduces optimization challenges, as reflected in the decreased DeAOT

performance with longer training clips (Table E, rows 1 and 3). Under such a setup and fair comparison, our full RMem has maintained comparable VOS quality compared with the baseline (rows 3 and 4). In conclusion, our RMem is also applicable for YoutubeVOS2019, although tuning the optimal hyper-parameters for training with longer sequence lengths is future work.

Index	Method	$\mathcal{G}$	$\mathcal{J}_s$	$\mathcal{J}_u$	$\mathcal{F}_s$	$\mathcal{F}_u$
1	DeAOT	85.9	84.6	89.4	80.8	88.9
2	DeAOT + RMem	85.9	84.6	89.4	80.8	88.9
3	DeAOT$^{\Psi}$	85.6	84.8	80.0	89.7	88.0
4	DeAOT$^{\Psi}$ + RMem	85.5	84.6	79.8	89.4	88.2

Table E. Analysis on YoutubeVOS2019 shows that, although not the primary focus of this paper, our RMem is also applicable for YoutubeVOS2019 with comparable performance with baselines. We first apply restricted memory banks to the original DeAOT checkpoint (rows 1 and 2). To enable temporal positional embedding (TPE), we train DeAOT under a longer sequence length and denote such models with "Ψ" (rows 3 and 4). The subscripts "s" and "u" denote the "seen" and "unseen" subsets of YoutubeVOS2019, respectively.

D. Additional Discussion on Limitations and Future Work

We briefly outlined the limitations of our study in Sec. 6 due to space limits. This section elaborates on more details.

As mentioned in Sec. 6, we prioritize the analysis of memory banks, and RMem is designed as a straightforward instantiation to demonstrate our insight. For this purpose, our study primarily engages with state-of-the-art methods like AOT [50] and DeAOT [48]. This choice is grounded, especially when common VOS studies are built upon a single or few preceding approaches due to the complexity of the framework, such as XMem [9], HODOR [1], and DeAOT [48]. One potential limitation could be that our RMem might implicitly depend on the transformer mechanisms and the affinity calculation in self-attention, which are adopted in AOT and DeAOT. These mechanisms natively support the temporal positional embedding and align with our key motivation of focusing the attention scores on relevant frames (Sec. 3 and Fig. 2). While future endeavors could explore adapting RMem for various VOS methods beyond the ones using transformers, near-future VOS methods will likely continue to employ a transformer-based framework, making our current RMem design compatible with them.

Another aspect mentioned in Sec. 6 is the potential for enhancing RMem with more advanced techniques. While the current simplicity of our approach effectively demonstrates our core insights into managing memory bank capacities, we acknowledge that it can benefit from a more sophisticated design. As especially pointed out in Sec. 6,

XMem [9] exhibits an intricate design for efficiently expanding memory banks. Though more complex than our current method of simply bounding memory bank sizes, such advancements could offer greater flexibility and potentially improve VOS.

Lastly, as discussed in Sec. 6, another option for enhancement lies in improving the decoding capabilities of the VOS framework. Our study maintains the original design of existing methods for a fair comparison, yet future research could explore scaling or modifying VOS architectures to further mitigate the challenges posed by expanding memory banks.