

Towards Temporally Consistent Referring Video Object Segmentation

Bo Miao, Mohammed Bennamoun, *Senior Member, IEEE*, Yongsheng Gao, *Senior Member, IEEE*,
Mubarak Shah, *Fellow, IEEE*, and Ajmal Mian, *Senior Member, IEEE*

Abstract—Referring Video Object Segmentation (R-VOS) methods face challenges in maintaining consistent object segmentation due to temporal context variability and the presence of other visually similar objects. We propose an *end-to-end* R-VOS paradigm that explicitly models temporal instance consistency alongside the referring segmentation. Specifically, we introduce a novel hybrid memory that facilitates inter-frame collaboration for robust spatio-temporal matching and propagation. Features of frames with automatically generated high-quality reference masks are propagated to segment the remaining frames based on multi-granularity association to achieve temporally consistent R-VOS. Furthermore, we propose a new Mask Consistency Score (MCS) metric to evaluate the temporal consistency of video segmentation. Extensive experiments demonstrate that our approach enhances temporal consistency by a significant margin, leading to top-ranked performance on popular R-VOS benchmarks, *i.e.*, Ref-YouTube-VOS (67.1%) and Ref-DAVIS17 (65.6%).

Index Terms—Referring video object segmentation, temporal consistency, deep learning.

I. INTRODUCTION

REFERRING video object segmentation (R-VOS) is a multimodal reasoning task, which aims at segmenting target objects in videos based on linguistic expressions. Unlike referring image segmentation [1]–[4], R-VOS poses a greater challenge as it requires maintaining temporal consistency and coherence throughout the entire video sequence. In contrast to semi-supervised video object segmentation [5]–[12] (Semi-VOS), R-VOS adds difficulty by not providing any ground-truth mask for the initial frame as reference.

Early R-VOS methods [13], [14] primarily employ convolution networks [15] for visual encoding and language grounding. However, their inherent limitations in capturing long-range dependencies and handling free-form features result in suboptimal performance. With the advancement in attention mechanism [16]–[20], various attention-based R-VOS approaches have emerged, leading to significant improvements. Techniques such as YOFO [21] and SgMg [22] employ cross-attention to model multimodal representations, enabling

flexible and long-range vision-language correspondence. Furthermore, ReferFormer [23] and MTTR [24] introduce end-to-end transformers with object queries to dynamically separate target objects from semantic features.

These transformer-based techniques employ a set of language-guided object queries to attend to multimodal features of different frames, eventually predicting conditional kernels [25] for each frame. The frame-specific conditional kernels are then used to generate object masks by convolving the corresponding multimodal features. Despite demonstrating promising performance, these approaches struggle to maintain consistent segmentation in complex and challenging scenarios. As illustrated in the first row of Fig. 1(b), the target sheep is accurately segmented in the initial frame, but the presence of other similar-looking sheep and changes in the local and global context, *e.g.*, the disappearance of the dog in the final frames, lead to the failure of the model to consistently segment the correct target.

To mitigate this inconsistency issue, some approaches utilize well-trained Semi-VOS models [26], [27] for post-processing, *e.g.*, ReferFormer [23] and PMINet [28] use CFBI [26] as a post-processing step to refine the masks. Such a naïve strategy requires a complex pipeline with additional backbones [15], [29] and training data [30]–[35], making R-VOS cumbersome and impractical. Moreover, this strategy does not yield satisfactory improvements in performance as we show in Table IV. An ideal solution would be an end-to-end paradigm that seamlessly integrates temporal consistency into R-VOS.

We introduce **Hybrid memory for Temporally consistent Referring video object segmentation (HTR)**. Instead of relying on well-trained Semi-VOS models, HTR integrates inter-frame association and referring segmentation into an end-to-end paradigm. This not only eliminates the need for additional backbones and training data, but also results in a significant gain in accuracy. In contrast to prior temporal propagation methods [26], [27], HTR includes a new hybrid memory that stores not only fine-grained local context but also lightweight global representations for explicit propagation. This memory exhibits higher robustness for feature propagation in the presence of noise that stems from imperfect automatically generated reference masks (as no ground-truth mask is provided for R-VOS during inference) and correlation between similar-looking objects. By leveraging inter-frame collaboration, highly confident segmentation in the hybrid memory is effectively propagated to the remaining frames, leading to improved temporal consistency and mask quality.

To evaluate the proposed HTR, we conduct a series of

Bo Miao, Mohammed Bennamoun, and Ajmal Mian are with the Department of Computer Science and Software Engineering, The University of Western Australia, Perth, Crawley, WA 6009 Australia.

Yongsheng Gao is with the School of Engineering, Griffith University, Brisbane, QLD 4111 Australia.

Mubarak Shah is with the Center for Research in Computer Vision, University of Central Florida, Orlando, FL 32816 USA.

This research was funded by the Australian Research Council Industrial Transformation Research Hub IH180100002. Professor Ajmal Mian is the recipient of an Australian Research Council Future Fellowship Award (project number FT210100268) funded by the Australian Government.

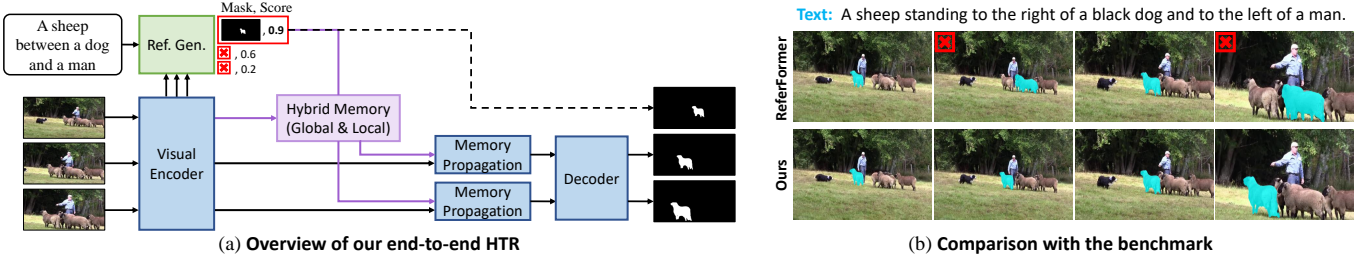


Fig. 1. (a) Our HTR automatically generates reference masks and achieves temporally consistent R-VOS in an end-to-end manner using the robust hybrid memory. (b) The baseline model ReferFormer [23] loses track of the target in some frames (marked by \times) whereas ours maintains temporal consistency in segmenting the correct object.

experiments on popular R-VOS benchmarks, Ref-YouTube-VOS [36], Ref-DAVIS17 [13], A2D-Sentences [14], JHMDB-Sentences [37]. Without bells and whistles, *e.g.*, test-time augmentations and model ensembles, HTR achieves top-ranked performance: 67.1% on Ref-YouTube-VOS, 65.6% on Ref-DAVIS17, 59.0% on A2D-Sentences, and 44.9% on JHMDB-Sentences. Notably, HTR outperforms the baseline model ReferFormer [23] by up to 5.1% points. Furthermore, we introduce a new metric, Mask Consistency Score (MCS), to evaluate the temporal consistency of video segmentation at different thresholds τ of individual frame segmentation accuracy. HTR shows significantly improved temporal consistency, with relative improvements of up to 53% in terms of MCS@0.9 ($\tau = 0.9$). Our main contributions are summarized below:

- We propose HTR, a highly effective and *end-to-end* paradigm that explicitly models temporal instance consistency alongside referring segmentation. HTR achieves top-ranked performance on the benchmark datasets while also having an efficient run-time speed.
- We develop a hybrid memory, in conjunction with inter-frame collaboration, to facilitate robust spatio-temporal propagation. The proposed memory significantly improves temporal consistency and mask quality.
- We propose a new Mask Consistency Score (MCS) metric to evaluate the temporal consistency of video segmentation. This will encourage progress in R-VOS by benchmarking future advancements in the temporal consistency of video segmentation.

II. RELATED WORK

A. Semi-supervised Video Object Segmentation

Semi-VOS aims to segment target objects throughout the entire video, given the ground-truth mask (of the first frame) [38]–[41]. Existing techniques mainly fall into three paradigms: finetuning-based, tracking-based, and matching-based. Finetuning-based methods dynamically update networks during test time using the first [42], [43], selected [44], or synthetic [45] frames from each video to make models object-specific. Although these methods show promising results, the slow test-time adaptation makes them inefficient. Tracking-based approaches [26], [46]–[55] avoid online learning and perform frame-to-frame propagation using masks of neighboring and annotated frames. However, they often lack

sufficient long-term context, leading to the loss of targets after occlusions.

To address the context limitations, matching-based methods [12], [56]–[63] construct memory banks with features from past frames and use attention mechanisms for temporal propagation. STM [64] and STCN [27] build memory banks to store mask and visual features at fixed intervals. SST [65] uses transformer blocks for affinity computation. AOT [11] and its subsequent works [66] establish hierarchical propagation to process multiple objects in parallel. XMem [10] introduces complementary memories inspired by the Atkinson-Shiffrin model [67] to efficiently handle long videos. In this work, we propose a hybrid memory that computes multi-granularity affinity for robust spatio-temporal propagation, even with imperfect reference masks.

B. Referring Video Object Segmentation

R-VOS methods conduct vision-and-language interactions to segment the referred object throughout the video [22], [23], [68]–[70]. LRE [13] performs referred object localization and regional segmentation. URVOS [36] combines vision-language and temporal features for iterative segmentation. MMTBVS [71] uses motion clues to enhance multimodal representations. MLRL [72] extracts video-level, frame-level, and object-level features for comprehensive representations. ClawcraneNet [73] and CMPC [74] mimic human behavior to progressively execute discriminative vision-language fusion. Similarly, HINet [75], YOFO [21], and LBDT [76] leverage multi-layer features for hierarchical cross-modal fusion.

Among these, transformer-based approaches have gained significant attention due to their end-to-end architecture and impressive performance. MTTR [24] introduces spatio-temporal transformer [77] for visual encoding and conditional kernels [25] for dynamic object segmentation. ReferFormer [23] employs language as queries to segment referred objects and includes an auxiliary detection loss to facilitate object localization learning. MANet [78] compresses video data and focuses on keyframe segmentation to speed up R-VOS. UNINEXT [79] introduces a unified model capable of training with data from various tasks. R²VOS [80] enhances multi-modal alignment through text reconstruction. OnlineRefer [81] proposes an online model with explicit query propagation. HTML [82] performs multimodal interactions across various temporal scales. SgMg [22] proposes a segment-and-optimize

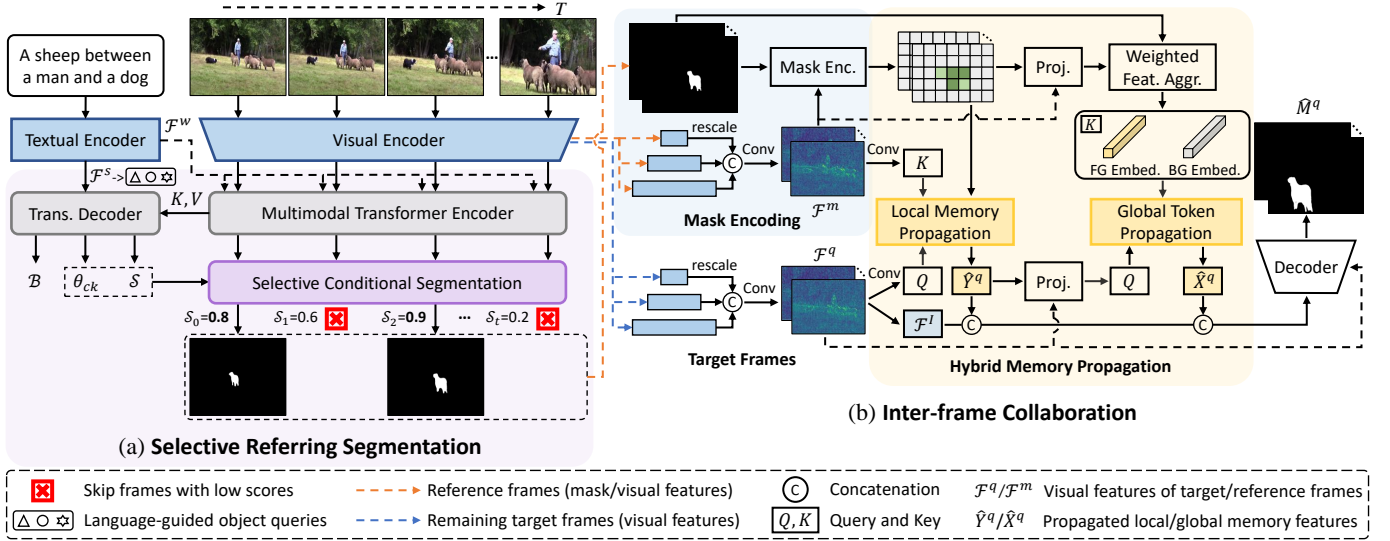


Fig. 2. Detailed architecture of HTR. (a) Selective referring process predicts the score S and conditional kernels θ_{ck} for each frame to selectively segments frames with high scores. Masks and visual features of these selected reference frames, and only the visual features of the remaining target frames are passed to (b) Inter-frame collaboration module. This module encodes reference frames to construct hybrid memory and aggregates memory features based on node-node (*local memory*) and node-object (*global token*) affinity to segment each pixel node in target frames. $\mathcal{F}^s/\mathcal{F}^w$: sentence/word features.

paradigm to solve the feature drift issue. This work proposes an end-to-end paradigm that seamlessly incorporates temporal consistency modeling to achieve consistent and accurate R-VOS.

III. METHOD

How can we devise an effective paradigm that accurately and consistently segments referred objects in videos? We answer this by proposing an end-to-end paradigm called HTR, which employs a novel memory module to propagate features of frames with high-quality masks to the remaining frames for achieving temporally consistent and accurate R-VOS, as shown in Fig. 2. HTR performs selective segmentation based on the confidence measure to generate high-quality reference masks for constructing the hybrid memory. The remaining frames with low scores are segmented by aggregating features from the memory based on robust inter-frame affinity. Compared to prior R-VOS methods [22]–[24], [78], [83], HTR can consistently segment the correct target across frames in an end-to-end manner, resulting in significant improvements in temporal consistency and segmentation quality.

A. Architecture

Figure 2 shows the architecture of our HTR, which consists of three components: feature encoding, selective referring segmentation, and inter-frame collaboration. The feature encoding stage extracts visual and textual features from the input frames and language expressions. The selective referring segmentation stage is built upon ReferFormer [23], which follows the standard transformer-based video segmentation framework [24], [84]. This stage employs language-guided object queries for conditional segmentation on frames with high scores. Notably, HTR is a general paradigm and is also compatible with other networks such as the recent SgMg [22]. The inter-frame

collaboration incorporates a new hybrid memory to perform robust inter-frame association and propagation for temporally consistent R-VOS.

B. Visual and Textual Encoders

We adopt Swin Transformer [85] or VideoSwin Transformer [77] as the visual encoder to extract visual features \mathcal{F}^v and RoBERTa [86] as the textual encoder for generating word-level features $\mathcal{F}^w \in \mathbb{R}^{N \times C}$ and sentence embeddings $\mathcal{F}^s \in \mathbb{R}^{1 \times C}$, owing to their strong representation ability. The extracted visual and linguistic features are projected to a lower dimension of 256 to facilitate efficient multimodal interactions. For each frame, the visual encoder produces a feature pyramid with spatial strides of $\{8, 16, 32\}$.

C. Selective Referring Segmentation

1) *Multimodal Transformer*: To generate object-related representations, cross-attention is employed to enhance each layer of the visual feature pyramid \mathcal{F}^v by incorporating linguistic information \mathcal{F}^w ,

$$\mathcal{F}_l^{vl} = \mathcal{F}_l^v \odot \text{Attn}(\mathcal{F}_l^v W_l^Q, \mathcal{F}^w W_l^K, \mathcal{F}^w W_l^V) \quad (1)$$

where \odot and W denote Hadamard product and learnable projection matrix. $\text{Attn}(Q, K, V)$ represents the standard attention block [16] with query, key, and value. The multimodal transformer integrates information from the vision-language feature pyramid \mathcal{F}^{vl} to effectively handle objects of varying sizes and scales. In this work, we adopt the memory-efficient Deformable Transformer [87] to encode vision-language features \mathcal{F}^{vl} and predict object embeddings \mathcal{E} , similar to [23], [88], [89]. The multimodal transformer encoder produces a feature pyramid of the same size as the input. The transformer decoder uses a set of language-guided object queries Q , generated by combining sentence features \mathcal{F}^s and learnable

embeddings, and cross-attends them with vision-language features \mathcal{F}^{vl} to predict frame-wise object embeddings \mathcal{E}

$$\mathcal{E} = \mathcal{Q} + \text{Attn}(\mathcal{Q}W^Q, \mathcal{F}^{vl}W^K, \mathcal{F}^{vl}W^V) \quad (2)$$

Each object embedding uses fully connected layers to predict conditional kernels θ_{ck} , a score \mathcal{S} , and a bounding box \mathcal{B} for an individual frame. The kernels θ_{ck} and score \mathcal{S} are employed for selective conditional segmentation, while the bounding box \mathcal{B} facilitates object localization learning following [23].

2) *Selective Conditional Segmentation*: Considering the temporal context variability, language-guided object queries and features may fail to consistently identify the correct targets, leading to temporally inconsistent and incorrect segmentation outcomes. To tackle this, we selectively apply conditional kernels θ_{ck} for referring segmentation on corresponding frames predicted to have high segmentation quality score \mathcal{S} . Specifically, we supervise the score \mathcal{S} during training based on the segmentation loss. The ground-truth score $\hat{\mathcal{S}}$ for the optimal object query with the minimum loss \mathcal{L} is set to one, while all others are set to zero:

$$\hat{\mathcal{S}}_i = \begin{cases} 1, & \text{if } \mathcal{L}_i = \min(\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_N) \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

Consequently, higher scores \mathcal{S} become indicative of superior segmentation quality. During inference, we rank the conditional kernels θ_{ck} based on their scores \mathcal{S} and selectively use the top-ranked kernels to segment their corresponding frames. For conditional segmentation, we use a cross-modal FPN decoder [23] to upsample the encoded vision-language features \mathcal{F}^{vl} of the selected frames. The corresponding kernels θ_{ck} are then employed as weights of point-wise convolutions applied to the decoded \mathcal{F}^{vl} to predict object masks \hat{M} . This selective segmentation stage efficiently produces reliable reference masks, facilitating subsequent inter-frame collaboration.

D. Inter-frame Collaboration

In our approach, frames with high scores are selectively segmented, and their features are propagated to segment the remaining target frames using the hybrid memory. As depicted in Fig. 2(b), the target frames aggregate memory features (\hat{X}^q and \hat{Y}^q) from the selected reference frames based on their inter-frame affinity. These aggregated features are then concatenated with the visual features of the target frames and passed to a convolutional decoder [90] to predict object masks.

1) *Mask Encoding*: HTR reuses multi-scale visual features for feature matching and enhancement. Unlike conventional memory networks [10], [27], [64] that use extra backbones for object mask encoding, HTR adopts an efficient fusion-based strategy. It directly projects 16×16 mask grids to hidden embeddings using a learnable transformation matrix and enhances these embeddings with visual features through two residual blocks [15] and a channel-spatial attention block [91].

E. Hybrid Memory

To effectively propagate spatio-temporal features in the presence of noise, which stems from imperfect reference

masks and correlation between similar-looking objects, we introduce a hybrid memory. As shown in Fig. 2(b), our memory comprises two complementary elements, local memory and global tokens, that collaboratively compute multi-granularity affinity for effective propagation.

1) *Local Memory*: The local memory captures pixel-level context for fine-grained feature matching and propagation, preserving object structure and appearance details. Let $\mathcal{F}^q \in \mathbb{R}^{HW \times C}$ denotes the visual features of a target frame, and $\mathcal{F}^m, Y^m \in \mathbb{R}^{T \times HW \times C}$ denote the memorized visual and mask features of the reference frames, respectively, where T, HW, C refer to the temporal, spatial, and channel dimensions. The propagation of local memory can be formulated as

$$\text{Sim}(Q, K)_{ij} = -\|Q_i - K_j\|_2^2 \quad (4)$$

$$\hat{Y}^q = \sigma(\text{Sim}(\mathcal{F}^q W^K, \mathcal{F}^m W^K)) Y^m \quad (5)$$

where $W^K \in \mathbb{R}^{C \times 64}$ is a trainable projection matrix used to extract the query and key for feature matching, σ denotes the Softmax operation applied along the spatio-temporal dimension of the local memory, and $\text{Sim}(Q, K)$ represents the L2 similarity [27] used to compute the affinity.

2) *Global Tokens*: The global tokens extract robust and lightweight representations for foreground and background, each condensed into a single token, providing crucial global spatio-temporal context. This facilitates target localization and alleviates error propagation. Specifically, based on the foreground and background probabilities ($M_{i=1}^m$ and $M_{i=0}^m$) obtained from selective referring segmentation, HTR performs weighted feature aggregation to extract foreground and background global representations. The aggregation process $\Phi_{aggr}(\mathcal{F}; M_i)$ is achieved through a weighted average of vision-mask joint representations from all reference frames,

$$\Phi_{aggr}(\mathcal{F}; M_i^m) = \frac{\sum_{j=1}^{T \times HW} (\mu(M_{i,j}^m - \tau) \odot M_{i,j}^m \odot \mathcal{F}_j)}{\sum_{j=1}^{T \times HW} (\mu(M_{i,j}^m - \tau) \odot M_{i,j}^m)} \quad (6)$$

where \odot denotes Hadamard product, $\tau = 0.5$ is the threshold, and $\mu(x)$ represents the Heaviside step function to extract foreground and background binary masks,

$$\mu(x) = \begin{cases} 1, & \text{if } x > 0, \\ 0, & \text{if } x \leq 0. \end{cases} \quad (7)$$

Through our weighted feature aggregation, noisy features with low scores contribute less to the global embeddings, resulting in robust foreground and background representations. The propagation of global tokens can be formulated as

$$\hat{X}^q = (\hat{Y}^q \parallel \mathcal{F}^q) W^J \Phi_{aggr}((Y^m \parallel \mathcal{F}^m) W^J; M^m)^T \quad (8)$$

where \parallel denotes the concatenation of visual and mask features and $W^J \in \mathbb{R}^{2C \times 64}$ is a trainable projection matrix to extract vision-mask joint representations. Note that we only compute the efficient node-object affinity using the global tokens to enhance the propagated fine-grained local memory features \hat{Y}^q . After hybrid memory propagation, the propagated features (\hat{X}^q and \hat{Y}^q) are concatenated with the visual features of the target frame. These features are then decoded using a common convolutional decoder [90] to predict the object mask \hat{M}^q .

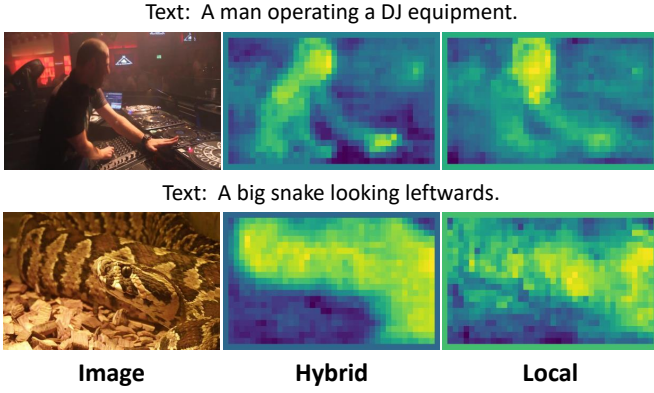


Fig. 3. Visualization of the propagated features using our hybrid memory and a standard local memory [27]. Our memory demonstrates robust feature propagation in challenging scenarios.

By leveraging our hybrid memory, HTR effectively handles spatio-temporal propagation (see Fig. 3) for temporally consistent R-VOS.

F. Mask Consistency Score

Temporal consistency in video segmentation is crucial as it ensures that object masks remain accurate and consistent throughout the video, regardless of changes in appearance or context. Inconsistent segmentation can cause flickering and loss of tracked targets. However, existing R-VOS models primarily use frame-wise accuracy metrics for evaluation. To the best of our knowledge, there are currently no dedicated metrics designed to evaluate the temporal consistency of R-VOS models.

To fill this gap, we propose a new metric, Mask Consistency Score (MCS), to evaluate the temporal consistency of video segmentation at specified thresholds. The segmentation of a video is considered consistent if the accuracy (Jaccard index) for *all* frames exceeds the threshold τ . Otherwise, the segmentation is deemed inconsistent due to the loss of tracking targets in one or more frames. The MCS metric is defined as the ratio of videos with consistent object segmentation. Let $\mathcal{J}_{x,y}$ denotes the Jaccard index of the segmentation for frame y in video x , the MCS at threshold τ is formulated as

$$\text{MCS@}\tau = \frac{1}{N} \sum_{x=1}^N \prod_{y=1}^T \mu(\mathcal{J}_{x,y} - \tau) \quad (9)$$

where N denotes the number of videos, T represents the number of frames per video, and $\mu(x)$ is the Heaviside step function mentioned in Eq. 7. In this work, we adopt thresholds of 0.1, 0.5, and 0.9 for evaluation. $\tau=0.1$ ensures at least some level of basic consistency (segmentation not completely off-target); $\tau=0.5$ ensures minimum acceptable temporal consistency for tracking; $\tau=0.9$ ensures high-quality results satisfying most real-world applications.

G. Training Loss Functions

The loss function of HTR consists of two parts:

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{refer}} + \mathcal{L}_{\text{prop}} \quad (10)$$

where $\mathcal{L}_{\text{refer}}$ and $\mathcal{L}_{\text{prop}}$ are losses for selective referring segmentation and inter-frame collaboration.

1) *Loss for Selective Referring Segmentation*: Selective segmentation predicts object masks \mathcal{M} , bounding boxes \mathcal{B} , and scores \mathcal{S} for each video. Given a set of predictions $y = \{\forall y_i, i \in [1, \dots, N]\}$ obtained from $N=5$ queries, where $y_i = \{\mathcal{M}_{i,j}, \mathcal{B}_{i,j}, \mathcal{S}_{i,j}\}_{j=1}^T$, we employ the Hungarian algorithm [92] to find the optimal assignment with the highest similarity to the ground truth for training. The selective referring segmentation loss can be represented as

$$\mathcal{L}_{\text{refer}} = \sum_{i=1}^N \lambda_{\mathcal{S}} \mathcal{L}_{\mathcal{S}_i} + \mathbf{1}_{y_i \neq \emptyset} [\lambda_{\mathcal{M}} \mathcal{L}_{\mathcal{M}_i} + \lambda_{\mathcal{B}} \mathcal{L}_{\mathcal{B}_i}] \quad (11)$$

where \mathcal{L} and λ denote the loss term and weight, $\mathbf{1}_{y_i \neq \emptyset}$ ensures that only the optimal query has mask and bounding box loss. Following [23], [84], [87], we use dice [93] and focal [94] loss for $\mathcal{L}_{\mathcal{M}}$, focal [94] loss for $\mathcal{L}_{\mathcal{S}}$, and GIoU [95] and L1 loss for $\mathcal{L}_{\mathcal{B}}$.

2) *Loss for Inter-frame Collaboration*: The masks predicted by inter-frame collaboration are supervised using bootstrapped cross entropy loss and dice loss with equal weighting, following [11], [27], [35].

IV. EXPERIMENTS

A. Datasets and Metrics

1) *Datasets*: We conduct experiments on four popular and challenging video datasets: Ref-YouTube-VOS [36], Ref-DAVIS17 [13], A2D-Sentences [14], JHMDB-Sentences [37]. Ref-YouTube-VOS is a large-scale dataset that consists of 3,471 training videos with 12,913 expressions and 507 validation videos with 2,096 expressions. Ref-DAVIS17 comprises 90 videos with 1,544 expressions, captured at 24 FPS. A2D-Sentences is an actor and action segmentation dataset with 3,782 video sequences, 6,656 action descriptions, and 1 to 5 frame annotations per sequence. JHMDB-Sentences includes 928 short video sequences and 928 descriptions covering 21 different action classes.

2) *Metrics*: We employ the standard metrics for evaluation. All results are evaluated using the official evaluation servers or codes. For Ref-YouTube-VOS and Ref-DAVIS17, Region Jaccard \mathcal{J} , Boundary Accuracy \mathcal{F} , and their average $\mathcal{J\&F}$ [99] over all videos are reported. In addition to accuracy metrics, we evaluate the temporal consistency of segmentation on Ref-YouTube-VOS using our proposed MCS metric. For A2D-Sentences and JHMDB-Sentences, we report the Precision@K, mAP over 0.50:0.05:0.95, Overall IoU, and Mean IoU. Precision@K measures the percentage of test samples with IoU scores exceeding a threshold K, where $K \in [0.5, 0.6, 0.7, 0.8, 0.9]$.

B. Implementation Details

Following [22]–[24], we train our models on the Ref-YouTube-VOS train split and evaluate them on both Ref-YouTube-VOS and Ref-DAVIS17. We also provide results of

TABLE I

QUANTITATIVE COMPARISON ON REF-YOUTUBE-VOS AND REF-DAVIS17. *: PRE-TRAINED WITH REFCOCO/g/+ IMAGES. S-HTR: INTEGRATED OUR HYBRID MEMORY WITH THE RECENT WORK SgMg [22] FOR COMPARISON. OUR MODELS ACHIEVE TOP-RANKED PERFORMANCE IN TERMS OF ACCURACY ($\mathcal{J}\&\mathcal{F}$) AND TEMPORALLY CONSISTENCY (MCS@ τ)

Method	Backbone	Ref-YouTube-VOS						Ref-DAVIS17		
		$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	MCS@0.1	MCS@0.5	MCS@0.9	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
CMSA [96]	ResNet-50	36.4	34.8	38.1	-	-	-	40.2	36.9	43.5
URVOS [36]	ResNet-50	47.2	45.3	49.2	-	-	-	51.5	47.3	56.0
CMPC-V [74]	I3D	47.5	45.6	49.3	-	-	-	-	-	-
PMINet [28]	ResNeSt-101	53.0	51.5	54.5	-	-	-	-	-	-
YOFO [21]	ResNet-50	48.6	47.5	49.7	-	-	-	53.3	48.8	57.8
LBDT [76]	ResNet-50	49.4	48.2	50.6	-	-	-	54.3	-	-
MLRL [72]	ResNet-50	49.7	48.4	51.0	-	-	-	52.8	50.0	55.4
UMVOS [97]	-	52.8	49.3	56.3	-	-	-	56.2	52.5	59.8
MTTR [24]	VideoSwin-T	55.3	54.0	56.6	50.6	39.0	7.1	-	-	-
MANet [78]	VideoSwin-T	55.6	54.8	56.5	-	-	-	-	-	-
ReferFormer [23]	VideoSwin-T	56.0	54.8	57.3	50.0	38.9	5.8	-	-	-
R ² VOS [80]	VideoSwin-T	57.1	55.9	58.2	-	-	-	-	-	-
SgMg [22]	VideoSwin-T	58.9	57.7	60.0	50.7	39.0	7.3	56.7	53.3	60.0
HTR (Ours)	VideoSwin-T	59.8	58.3	61.3	60.4	47.8	8.2	57.2	53.8	60.6
S-HTR (Ours)	VideoSwin-T	61.5	60.0	63.0	62.2	50.5	8.5	58.5	55.9	61.0
<i>Additionally pre-trained with RefCOCO/g/+ static images</i>										
ReferFormer* [23]	VideoSwin-T	59.4	58.0	60.9	50.5	40.8	5.9	59.5	56.3	62.7
HTML* [82]	VideoSwin-T	61.2	59.5	63.0	-	-	-	-	-	-
R ² VOS* [80]	VideoSwin-T	61.3	59.6	63.1	-	-	-	-	-	-
HTR* (Ours)	VideoSwin-T	64.2	62.7	65.7	63.3	53.8	10.9	63.2	59.4	67.0
ReferFormer* [23]	VideoSwin-B	62.9	61.3	64.6	54.8	46.3	8.0	61.1	58.1	64.1
OnlineRefer* [81]	VideoSwin-B	62.9	61.0	64.7	54.4	43.2	7.1	62.4	59.1	65.6
HTML* [82]	VideoSwin-B	63.4	61.5	65.2	-	-	-	62.1	59.2	65.1
VLT* [2]	VideoSwin-B	63.8	61.9	65.6	-	-	-	61.6	58.9	64.3
SgMg* [22]	VideoSwin-B	65.7	63.9	67.4	55.7	48.4	11.4	63.3	60.6	66.0
HTR* (Ours)	VideoSwin-B	66.7	64.9	68.6	64.8	56.7	12.2	63.5	60.5	66.6
S-HTR* (Ours)	VideoSwin-B	67.5	65.6	69.4	65.0	56.2	12.8	64.8	61.8	67.9
ReferFormer* [23]	Swin-L	62.4	60.8	64.0	53.2	44.7	7.6	60.5	57.6	63.4
OnlineRefer* [81]	Swin-L	63.5	61.6	65.5	54.0	44.8	7.7	64.8	61.6	67.7
DMFormer* [98]	Swin-L	64.9	63.4	66.5	-	-	-	62.3	59.5	65.1
HTR* (Ours)	Swin-L	67.1	65.3	68.9	64.0	54.4	11.6	65.6	62.3	68.8

our models that are pre-trained on static image datasets Ref-COCO/g/+ [100], [101] and then fine-tuned on Ref-YouTube-VOS. For A2D-Sentences and JHMDB-Sentences, we fine-tune the pre-trained models on the A2D-Sentences train split and evaluate them on both A2D-Sentences and JHMDB-Sentences.

We maintain the same optimization strategies as [22], [23] for a fair comparison. Our models are optimized using the AdamW [102] optimizer with a weight decay of 5×10^{-4} . The initial learning rates are set to 2.5×10^{-5} for the visual encoder and inter-frame collaboration module, 5×10^{-6} for the textual encoder, and 5×10^{-5} for the remaining components. The batch size is 4/2 for pre-training/main training and distributed on two 24GB RTX 3090 GPUs. During pre-training, each sample consists of one static image, we train our models for 12 epochs with a learning rate decay by a factor of 10 at the 8th and 10th epochs. During main training, each sample consists of 5 randomly sampled frames, we freeze the text encoder and train the models for 6/10 epochs with a learning rate decay by a factor of 10 at the 3/6th and 5/8th epochs, depending on whether pre-training was performed. The loss function coefficients, λ_{dice} , λ_{L1} , λ_{focal} , λ_{GIoU} , λ_{ce} are set to 5, 5, 2, 2, 1, respectively. During inference, we follow [22], [23] to perform video-level segmentation (N =frame number)

on Ref-YouTube-VOS and per-clip segmentation ($N=36$) on Ref-DAVIS17. We selectively segment the top 25% of ranked frames to construct memory for each video (see Table VI for the ablation study).

C. Main Results

1) *Results on Ref-YouTube-VOS*: Table I shows that our HTR (with Swin-L) achieves the top-ranked performance of **67.1** $\mathcal{J}\&\mathcal{F}$ on Ref-YouTube-VOS, significantly outperforming the baseline ReferFormer [23] by **4.7% points**. Without pre-training on Ref-COCO/g/+, HTR with VideoSwin-T achieves 59.8 $\mathcal{J}\&\mathcal{F}$ at 58 FPS, still surpassing ReferFormer while running about $1.2 \times$ faster. Despite its selective segmentation stage being built upon ReferFormer, HTR exceeds all recent competitors, including OnlineRefer [81], HTML [82], and SgMg [22], particularly in terms of temporal consistency. This demonstrates the effectiveness of our approach. Additionally, we integrate our hybrid memory with the recent work SgMg [22] to create the end-to-end model S-HTR for comparison. This integration further improves performance by a large margin. Without pre-training on Ref-COCO/g/+, S-HTR with VideoSwin-T achieves 61.5 $\mathcal{J}\&\mathcal{F}$, surpassing SgMg by 2.6% points. These results further demonstrate our architecture is effective and generic.

TABLE II
QUANTITATIVE COMPARISON TO STATE-OF-THE-ART R-VOS METHODS ON A2D-SENTENCES

		mAP	IoU		Precision				
Method	Backbone		Overall	Mean	P@0.5	P@0.6	P@0.7	P@0.8	P@0.9
Hu <i>et al.</i> [14]	VGG-16	13.2	47.4	35.0	34.8	23.6	13.3	3.3	0.1
Gavrilyuk <i>et al.</i> [103]	I3D	19.8	53.6	42.1	47.5	34.7	21.1	8.0	0.2
CMSA + CFSA [96]	ResNet-101	-	61.8	43.2	48.7	43.1	35.8	23.1	5.2
ACAN [104]	I3D	27.4	60.1	49.0	55.7	45.9	31.9	16.0	2.0
CMPC-V [74]	I3D	40.4	65.3	57.3	65.5	59.2	50.6	34.2	9.8
ClawCraneNet [73]	ResNet-50/101	-	63.1	59.9	70.4	67.7	61.7	48.9	17.1
MTTR [24]	VideoSwin-T	46.1	72.0	64.0	75.4	71.2	63.8	48.5	16.9
ReferFormer [23]	VideoSwin-B	55.0	78.6	70.3	83.1	80.4	74.1	57.9	21.2
OnlineRefer [81]	VideoSwin-B	-	79.6	70.5	83.1	80.2	73.4	56.8	21.7
HTML [82]	VideoSwin-B	56.7	79.5	71.2	84.0	81.5	75.8	59.2	22.8
HTR (Ours)	VideoSwin-B	59.0	80.1	72.6	85.2	83.2	77.9	62.4	25.5

TABLE III
QUANTITATIVE COMPARISON TO STATE-OF-THE-ART R-VOS METHODS ON JHMDB-SENTENCES

		mAP	IoU		Precision				
Method	Backbone		Overall	Mean	P@0.5	P@0.6	P@0.7	P@0.8	P@0.9
Hu <i>et al.</i> [14]	VGG-16	17.8	54.6	52.8	63.3	35.0	8.5	0.2	0.0
Gavrilyuk <i>et al.</i> [103]	I3D	23.3	54.1	54.2	69.9	46.0	17.3	1.4	0.0
CMSA + CFSA [96]	ResNet-101	-	62.8	58.1	76.4	62.5	38.9	9.0	0.1
ACAN [104]	I3D	28.9	57.6	58.4	75.6	56.4	28.7	3.4	0.0
CMPC-V [74]	I3D	34.2	61.6	61.7	81.3	65.7	37.1	7.0	0.0
ClawCraneNet [73]	ResNet-50/101	-	64.4	65.6	88.0	79.6	56.6	14.7	0.2
MTTR [24]	VideoSwin-T	39.2	70.1	69.8	93.9	85.2	61.6	16.6	0.1
ReferFormer [23]	VideoSwin-B	43.7	73.0	71.8	96.2	90.2	70.2	21.0	0.3
OnlineRefer [81]	VideoSwin-B	-	73.5	71.9	96.1	90.4	71.0	21.9	0.2
HTR (Ours)	VideoSwin-B	44.9	73.9	72.6	97.1	92.0	72.1	23.7	0.3

2) *Results on Ref-DAVIS17*: In Table I, HTR shows outstanding generality on Ref-DAVIS17, remarkably surpassing the benchmark ReferFormer [23] by **5.1% points** in terms of $\mathcal{J}\&\mathcal{F}$. Specifically, HTR with VideoSwin-T achieves 63.2 $\mathcal{J}\&\mathcal{F}$, while HTR with Swin-L achieves the top-ranked performance at **65.6 $\mathcal{J}\&\mathcal{F}$** .

3) *Temporal Consistency*: In Table I, we compare HTR with other open-sourced R-VOS models on Ref-YouTube-VOS using the proposed MCS metric, which evaluates the temporal consistency of segmentation across thresholds from 0.1 to 0.9. HTR significantly outperforms ReferFormer [23], MTTR [24], OnlineRefer [81], and SgMg [22] across all thresholds due to the temporal consistency modeling. When Swin-L is used as the visual encoder, HTR outperforms the baseline ReferFormer by 20.3%, 21.7%, and 52.6% under MCS@0.1, MCS@0.5, MCS@0.9, respectively. Although HTR with VideoSwin-T employs a smaller backbone compared to ReferFormer with Swin-L, it still achieves notably higher MCS. These results prove the effectiveness of HTR in improving temporal consistency. Moreover, the MCS decreases remarkably for thresholds above 0.5, indicating the future research direction of developing R-VOS models that can produce both temporally consistent and accurate segmentation outcomes.

4) *Results on A2D-Sentences & JHMDB-Sentences*: In Table II and Table III, we further evaluate our models on the A2D-Sentences and JHMDB-Sentences, comparing them with other state-of-the-art methods. As only a few frames in each video sequence are annotated for training, the models

are less impacted by temporally consistent modeling on these two datasets. Despite this, our models exceed other R-VOS methods in all metrics, further demonstrating the effectiveness of our approach.

5) *Inference Speed and Parameters*: We perform selective conditional segmentation and efficient inter-frame collaboration for temporally consistent R-VOS. By performing conditional segmentation *only* on top-ranked (reference) frames, *i.e.*, avoiding the slow and computationally heavy cross-modal FPN decoding on the remaining frames, HTR runs 8-16% faster than ReferFormer [23], despite having about 12% (22M) more parameters attributed to the hybrid memory. This demonstrates the great performance trade-off in our approach.

D. Qualitative Results

We visualize some typical segmentation results in Fig. 4. ReferFormer and SgMg face challenges when dealing with temporal context changes, such as the target giraffe is not always in the middle, and the existence of other similar-looking objects like giraffes and gorillas. In contrast, HTR predicts correct and temporally consistent object masks, thanks to its temporal consistency modeling. Additionally, SgMg can handle various challenging scenarios including fast motion, occlusion, objects with similar appearance, and small objects, as shown in Fig. 4(b).



Fig. 4. Qualitative results. (a) Our HTR predicts temporally consistent results compared to ReferFormer [23] and SgMg [22]. (b) HTR can handle various challenging scenarios, including fast motion, objects with similar appearance, occlusion, and small objects. (c) HTR fails to improve mask quality without good reference masks.

E. Limitations

Our paradigm relies upon the availability of a few high-quality reference masks. However, the selective segmentation stage may not provide any good reference mask for some challenging videos. One such failure case is shown in Fig. 4(c). In such cases, our paradigm will obviously be unable to improve the segmentation quality. Nevertheless, our approach still significantly outperforms existing methods without using

any ground-truth reference mask. Future research could develop more effective selective referring segmentation models.

F. Visualization of Global Token Association

In Fig. 5, we visualize the affinity maps \hat{X}^q between the vision-mask joint representations $F^J \in \mathbb{R}^{C \times HW}$ of target frames and the global foreground representation $E^m \in \mathbb{R}^{C \times 1}$ in the hybrid memory (as defined in Eq. 8). Our global tokens

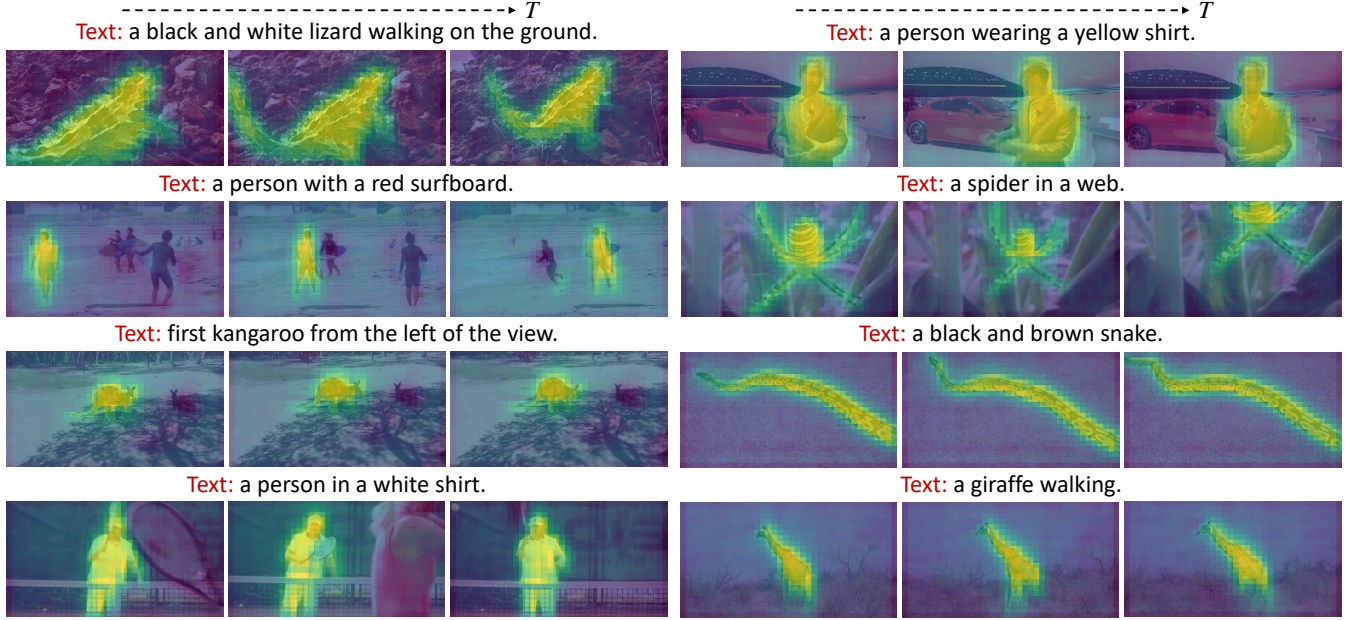


Fig. 5. Visualization of the affinity between the vision-mask joint representations of target frames and the global foreground representation in the hybrid memory. Our feature aggregation extracts robust global representations to localize the correct targets.

TABLE IV

COMPARISON WITH WELL-TRAINED VOS MODELS ON REF-YOUTUBE-VOS. E.: END-TO-END. ADD. ENC.: ADDITIONAL VISUAL ENCODER. ADD. DATA: ADDITIONAL TRAINING DATA. X-HTR: REPLACE OUR HYBRID MEMORY WITH XMEM [10]. HTR IS END-TO-END AND ACHIEVES SUPERIOR PERFORMANCE UNDER FAIR COMPARISON

Methods	E.	Add. Enc.	Add. Data	$\mathcal{J}\&\mathcal{F}$	Prop. FPS
ReferFormer	✓	×	×	62.4	-
ReferFormer + CFBI	×	✓	✓	63.3	7
ReferFormer + STCN	×	✓	✓	63.6	69
ReferFormer + XMem	×	✓	✓	64.4	61
X-HTR (Ours)	✓	×	×	65.0	168
HTR (Ours)	✓	×	×	67.1	181

successfully localize the referred objects thanks to the robust global token aggregation.

G. Ablation Study

1) *Well-trained VOS models*: In Table IV, we compare HTR with the combination of the benchmark ReferFormer [23] and well-trained temporal propagation models, *i.e.*, CFBI [26], STCN [27], and XMem [10]. Following ReferFormer, we use Swin-L as the visual encoder in this experiment (VideoSwin-T is used in subsequent ablations). We perform comparisons using 360p frames and adopt high-score frames as reference frames for fair comparison. Our end-to-end HTR significantly surpasses others by over 2.7% and runs up to 26× faster. Moreover, replacing the hybrid memory with XMem in our paradigm results in a 2.1% lower $\mathcal{J}\&\mathcal{F}$. XMem designs sensory memory to propagate the ground-truth mask of the first frame which does not align with our challenge of propagating imperfect reference masks from any frames. In contrast, HTR extracts robust global context and improves matching and memory extracted from imperfect reference masks. These

TABLE V

ABLATION OF MEMORY COMPONENTS. COLLABORATIVE LOCAL AND GLOBAL COMPONENTS ACHIEVE THE BEST PERFORMANCE

Methods	$\mathcal{J}\&\mathcal{F}$	MCS@0.1	MCS@0.5
w/o memory	59.4	50.5	40.8
w/ Local	62.5	60.8	50.7
w/ Global	62.9	63.3	53.0
w/ Hybrid	64.2	63.3	53.8

TABLE VI

ABLATION OF MEMORIZED RATIO

Mem. Ratio (k)	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
1%	63.5	62.0	64.9
25%	64.2	62.7	65.7
50%	63.9	62.5	65.3
75%	63.2	61.7	64.6

results confirm the effectiveness of our hybrid memory and paradigm.

2) *Hybrid Memory*: In Table V, we evaluate the individual contributions of the components within our hybrid memory. Both local and global components exhibit strong temporal propagation ability in our paradigm. Notably, global tokens outperform local memory. Our global tokens employ weighted feature aggregation to suppress noise and form robust global representations. This is crucial for improving the temporal consistency and mask quality during propagation especially when the initial reference masks are automatically generated and imperfect. Overall, the collaboration of local and global components leads to superior performance.

3) *Memorized Frames*: In Table VI, we evaluate the performance of HTR across different ratios of memorized frames. HTR dynamically selects the top k percent of ranked frames

TABLE VII
ABLATION OF STRIDES

Stride	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
8	63.4	62.1	64.8
16	64.2	62.7	65.7
32	63.3	61.8	64.8

TABLE VIII
INTEGRATE HTR TO LARGE R-VOS MODEL

Method	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
UNINEXT [79]	70.1	67.5	72.6
Ours	72.5	70.2	74.8

#	User	Entries	Date of Last Entry	Overall \uparrow	$\mathcal{J} \uparrow$	$\mathcal{F} \uparrow$
1	GTR.	2	05/18/23	0.725 (1)	0.702 (2)	0.748 (1)
2	JerryX	10	05/22/22	0.722 (2)	0.703 (1)	0.742 (3)
3	noff	219	05/25/22	0.719 (3)	0.694 (3)	0.745 (2)
4	dian	453	05/07/23	0.706 (4)	0.683 (4)	0.728 (4)
5	MasterBin-IIAU	35	03/08/23	0.701 (5)	0.676 (5)	0.727 (5)

Fig. 6. Our model (GTR. in the figure) is the **top performer** on the active 4th R-VOS competition leaderboard, captured in May 2023.

(frame numbers are rounded up) to construct the hybrid memory. The results show that memorizing the top 25% of ranked frames to segment the remaining 75% achieves the best performance.

4) *Different Feature Strides*: In this work, HTR performs inter-frame collaboration using features with stride 16, a common setting for memory modules [10], [64]. Table VII presents our evaluation of different strides, demonstrating that the common setting achieves the best performance.

H. Collaboration with Large R-VOS Model

Our trained HTR models can seamlessly collaborate with other large R-VOS models to boost temporal consistency and performance even without re-training since inter-frame collaboration takes general masks as input. We leverage UNINEXT [79] to generate reference masks for our trained model. We refer to this variant as UNI-HTR for brevity. UNINEXT is a recent large model that achieves impressive performance. It uses the ViT-Huge [105] visual encoder and 32 A100 GPUs to train on 11 large-scale training datasets from related tasks. In Table VIII, our approach improves the performance by 2.4% on Ref-YouTube-VOS *without* techniques like multi-scale inference and model ensembles. This leads to top position on the active leaderboard in the 4th R-VOS competition on CodaLab, as evidenced in Fig. 6 (captured in May 2023).

V. CONCLUSION

This paper proposes a novel end-to-end paradigm for referring video object segmentation (R-VOS) that effectively addresses the temporal inconsistency issue. Unlike previous solutions, our paradigm associates and propagates spatio-temporal features in an end-to-end fashion for achieving temporally consistent segmentation. To handle spatio-temporal

propagation in the presence of noise, we introduced the hybrid memory with inter-frame collaboration to conduct robust multi-granularity association. Furthermore, we proposed a new Mask Consistency Score (MCS) metric to evaluate the temporal consistency of video segmentation. Our results demonstrate that incorporating explicit temporal instance consistency modeling not only effectively improves temporal consistency, but also greatly enhances segmentation quality, achieving top-ranked performance on popular R-VOS benchmarks. We hope our paradigm will also benefit other video segmentation tasks in the future.

REFERENCES

- [1] Z. Yang, J. Wang, Y. Tang, K. Chen, H. Zhao, and P. H. Torr, “Lavt: Language-aware vision transformer for referring image segmentation,” in *CVPR*, 2022, pp. 18 155–18 165.
- [2] H. Ding, C. Liu, S. Wang, and X. Jiang, “Vlt: Vision-language transformer and query generation for referring segmentation,” *TPAMI*, vol. 45, no. 06, pp. 7900–7916, 2023.
- [3] C. Shang, H. Li, H. Qiu, Q. Wu, F. Meng, T. Zhao, and K. N. Ngan, “Cross-modal recurrent semantic comprehension for referring image segmentation,” *IEEE Trans. Circuits Syst. Video Technol.*, 2022.
- [4] H. Li, M. Sun, J. Xiao, E. G. Lim, and Y. Zhao, “Fully and weakly supervised referring expression segmentation with end-to-end learning,” *IEEE Trans. Circuits Syst. Video Technol.*, 2023.
- [5] J. Fan, B. Liu, K. Zhang, and Q. Liu, “Semi-supervised video object segmentation via learning object-aware global-local correspondence,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 12, pp. 8153–8164, 2021.
- [6] W. Zhu, J. Li, J. Lu, and J. Zhou, “Separable structure modeling for semi-supervised video object segmentation,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 1, pp. 330–344, 2021.
- [7] W. Liu, G. Lin, T. Zhang, and Z. Liu, “Guided co-segmentation network for fast video object segmentation,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 4, pp. 1607–1617, 2020.
- [8] Z. Tan, B. Liu, Q. Chu, H. Zhong, Y. Wu, W. Li, and N. Yu, “Real time video object segmentation in compressed domain,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 31, no. 1, pp. 175–188, 2020.
- [9] F. Lin, H. Xie, C. Liu, and Y. Zhang, “Bilateral temporal re-aggregation for weakly-supervised video object segmentation,” *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 7, pp. 4498–4512, 2021.
- [10] H. K. Cheng and A. G. Schwing, “Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model,” in *ECCV*, 2022, pp. 640–658.
- [11] Z. Yang, Y. Wei, and Y. Yang, “Associating objects with transformers for video object segmentation,” in *NeurIPS*, 2021, pp. 2491–2502.
- [12] B. Miao, M. Bennamoun, Y. Gao, and A. Mian, “Region aware video object segmentation with deep motion modeling,” *arXiv preprint arXiv:2207.10258*, 2022.
- [13] A. Khoreva, A. Rohrbach, and B. Schiele, “Video object segmentation with language referring expressions,” in *ACCV*, 2019, pp. 123–141.
- [14] R. Hu, M. Rohrbach, and T. Darrell, “Segmentation from natural language expressions,” in *ECCV*, 2016, pp. 108–124.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016, pp. 770–778.
- [16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *NeurIPS*, 2017.
- [17] B. Miao, L. Zhou, A. S. Mian, T. L. Lam, and Y. Xu, “Object-to-scene: Learning to transfer object knowledge to indoor scene recognition,” in *IROS*. IEEE, 2021, pp. 2069–2075.
- [18] M. Feng, H. Hou, L. Zhang, Y. Guo, H. Yu, Y. Wang, and A. Mian, “Exploring hierarchical spatial layout cues for 3d point cloud based scene graph prediction,” *IEEE Transactions on Multimedia*, 2023.
- [19] Z. Wu, M. Feng, Y. Wang, H. Xie, W. Dong, B. Miao, and A. Mian, “External knowledge enhanced 3d scene generation from sketch,” *arXiv preprint arXiv:2403.14121*, 2024.
- [20] Z. Wu, Y. Wang, M. Feng, H. Xie, and A. Mian, “Sketch and text guided diffusion model for colored point cloud generation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8929–8939.

- [21] D. Li, R. Li, L. Wang, Y. Wang, J. Qi, L. Zhang, T. Liu, Q. Xu, and H. Lu, "You only infer once: Cross-modal meta-transfer for referring video object segmentation," in *AAAI*, 2022, pp. 1297–1305.
- [22] B. Miao, M. Bennamoun, Y. Gao, and A. Mian, "Spectrum-guided multi-granularity referring video object segmentation," in *ICCV*, 2023, pp. 920–930.
- [23] J. Wu, Y. Jiang, P. Sun, Z. Yuan, and P. Luo, "Language as queries for referring video object segmentation," in *CVPR*, 2022, pp. 4974–4984.
- [24] A. Botach, E. Zheltonozhskii, and C. Baskin, "End-to-end referring video object segmentation with multimodal transformers," in *CVPR*, 2022, pp. 4985–4995.
- [25] Z. Tian, C. Shen, and H. Chen, "Conditional convolutions for instance segmentation," in *ECCV*, 2020, pp. 282–298.
- [26] Z. Yang, Y. Wei, and Y. Yang, "Collaborative video object segmentation by foreground-background integration," in *ECCV*, 2020, pp. 332–348.
- [27] H. K. Cheng, Y.-W. Tai, and C.-K. Tang, "Rethinking space-time networks with improved memory coverage for efficient video object segmentation," in *NeurIPS*, 2021, pp. 11 781–11 794.
- [28] Z. Ding, T. Hui, S. Huang, S. Liu, X. Luo, J. Huang, and X. Wei, "Progressive multimodal interaction network for referring video object segmentation," *The 3rd Large-scale Video Object Segmentation Challenge*, p. 7, 2021.
- [29] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *ECCV*, 2018, pp. 801–818.
- [30] L. Wang, H. Lu, Y. Wang, M. Feng, D. Wang, B. Yin, and X. Ruan, "Learning to detect salient objects with image-level supervision," in *CVPR*, 2017, pp. 136–145.
- [31] J. Shi, Q. Yan, L. Xu, and J. Jia, "Hierarchical image saliency detection on extended cssd," *TPAMI*, vol. 38, no. 4, pp. 717–729, 2015.
- [32] Y. Zeng, P. Zhang, J. Zhang, Z. Lin, and H. Lu, "Towards high-resolution salient object detection," in *ICCV*, 2019, pp. 7234–7243.
- [33] H. K. Cheng, J. Chung, Y.-W. Tai, and C.-K. Tang, "Cascadepsp: Toward class-agnostic and very high-resolution segmentation via global and local refinement," in *CVPR*, 2020, pp. 8890–8899.
- [34] X. Li, T. Wei, Y. P. Chen, Y.-W. Tai, and C.-K. Tang, "Fss-1000: A 1000-class dataset for few-shot segmentation," in *CVPR*, 2020, pp. 2869–2878.
- [35] H. K. Cheng, Y.-W. Tai, and C.-K. Tang, "Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion," in *CVPR*, 2021, pp. 5559–5568.
- [36] S. Seo, J.-Y. Lee, and B. Han, "Urvos: Unified referring video object segmentation network with a large-scale benchmark," in *ECCV*, 2020, pp. 208–223.
- [37] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black, "Towards understanding action recognition," in *ICCV*, 2013, pp. 3192–3199.
- [38] Y. Lu, J. Zhang, S. Sun, Q. Guo, Z. Cao, S. Fei, B. Yang, and Y. Chen, "Label-efficient video object segmentation with motion clues," *IEEE Trans. Circuits Syst. Video Technol.*, 2023.
- [39] X. Xu, J. Zhao, J. Wu, and F. Shen, "Switch and refine: A long-term tracking and segmentation framework," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 33, no. 3, pp. 1291–1304, 2022.
- [40] Y. Chen, D. Zhang, Y. Zheng, Z.-X. Yang, E. Wu, and H. Zhao, "Boosting video object segmentation via robust and efficient memory network," *IEEE Trans. Circuits Syst. Video Technol.*, 2023.
- [41] Y. Gui, Y. Tian, D.-J. Zeng, Z.-F. Xie, and Y.-Y. Cai, "Reliable and dynamic appearance modeling and label consistency enforcing for fast and coherent video object segmentation with the bilateral grid," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 12, pp. 4781–4795, 2019.
- [42] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool, "One-shot video object segmentation," in *CVPR*, 2017, pp. 221–230.
- [43] K.-K. Maninis, S. Caelles, Y. Chen, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool, "Video object segmentation without temporal information," *TPAMI*, vol. 41, no. 6, pp. 1515–1530, 2018.
- [44] T. Meinhardt and L. Leal-Taixé, "Make one-shot video object segmentation efficient again," in *NeurIPS*, 2020, pp. 10 607–10 619.
- [45] J. Luiten, P. Voigtlaender, and B. Leibe, "Premvos: Proposal-generation, refinement and merging for video object segmentation," in *ACCV*, 2018, pp. 565–580.
- [46] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung, "Learning video object segmentation from static images," in *CVPR*, 2017, pp. 2663–2672.
- [47] S. W. Oh, J.-Y. Lee, K. Sunkavalli, and S. J. Kim, "Fast video object segmentation by reference-guided mask propagation," in *CVPR*, 2018, pp. 7376–7385.
- [48] L. Yang, Y. Wang, X. Xiong, J. Yang, and A. K. Katsaggelos, "Efficient video object segmentation via network modulation," in *CVPR*, 2018, pp. 6499–6507.
- [49] X. Chen, Z. Li, Y. Yuan, G. Yu, J. Shen, and D. Qi, "State-aware tracker for real-time video object segmentation," in *CVPR*, 2020, pp. 9384–9393.
- [50] J. Cheng, Y.-H. Tsai, W.-C. Hung, S. Wang, and M.-H. Yang, "Fast and accurate online video object segmentation via tracking parts," in *CVPR*, 2018, pp. 7415–7424.
- [51] H. Lin, X. Qi, and J. Jia, "Agss-vos: Attention guided single-shot video object segmentation," in *ICCV*, 2019, pp. 3949–3957.
- [52] S. Xu, D. Liu, L. Bao, W. Liu, and P. Zhou, "Mhp-vos: Multiple hypotheses propagation for video object segmentation," in *CVPR*, 2019, pp. 314–323.
- [53] L. Zhang, Z. Lin, J. Zhang, H. Lu, and Y. He, "Fast video object segmentation via dynamic targeting network," in *ICCV*, 2019, pp. 5582–5591.
- [54] P. Voigtlaender, Y. Chai, F. Schroff, H. Adam, B. Leibe, and L.-C. Chen, "Feelvos: Fast end-to-end embedding learning for video object segmentation," in *CVPR*, 2019, pp. 9481–9490.
- [55] Z. Wang, J. Xu, L. Liu, F. Zhu, and L. Shao, "Ranet: Ranking attention network for fast video object segmentation," in *ICCV*, 2019, pp. 3978–3987.
- [56] H. Seong, J. Hyun, and E. Kim, "Kernelized memory network for video object segmentation," in *ECCV*, 2020, pp. 629–645.
- [57] B. Miao, M. Bennamoun, Y. Gao, and A. Mian, "Self-supervised video object segmentation by motion-aware mask propagation," in *ICME*, 2022, pp. 1–6.
- [58] K. Duarte, Y. S. Rawat, and M. Shah, "Capsulevos: Semi-supervised video object segmentation using capsule routing," in *ICCV*, 2019, pp. 8480–8489.
- [59] Y. Zhang, Z. Wu, H. Peng, and S. Lin, "A transductive approach for video object segmentation," in *CVPR*, 2020, pp. 6949–6958.
- [60] X. Xu, J. Wang, X. Li, and Y. Lu, "Reliable propagation-correction modulation for video object segmentation," in *AAAI*, 2022, pp. 2946–2954.
- [61] L. Hu, P. Zhang, B. Zhang, P. Pan, Y. Xu, and R. Jin, "Learning position and target consistency for memory-based video object segmentation," in *CVPR*, 2021, pp. 4144–4154.
- [62] Y. Liang, X. Li, N. Jafari, and J. Chen, "Video object segmentation with adaptive feature bank and uncertain-region refinement," in *NeurIPS*, 2020, pp. 3430–3441.
- [63] B. Miao, M. Bennamoun, Y. Gao, and A. Mian, "Regional video object segmentation by efficient motion-aware mask propagation," in *DICTA*, 2022, pp. 1–6.
- [64] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim, "Video object segmentation using space-time memory networks," in *ICCV*, 2019, pp. 9226–9235.
- [65] B. Duke, A. Ahmed, C. Wolf, P. Aarabi, and G. W. Taylor, "Sstvos: Sparse spatiotemporal transformers for video object segmentation," in *CVPR*, 2021, pp. 5912–5921.
- [66] Z. Yang and Y. Yang, "Decoupling features in hierarchical propagation for video object segmentation," *NeurIPS*, pp. 36 324–36 336, 2022.
- [67] R. C. Atkinson and R. M. Shiffrin, "Human memory: A proposed system and its control processes," in *Psychol. Learn. Motiv.*, 1968, vol. 2, pp. 89–195.
- [68] J. Tang, G. Zheng, and S. Yang, "Temporal collection and distribution for referring video object segmentation," in *ICCV*, 2023, pp. 15 466–15 476.
- [69] Z. Luo, Y. Xiao, Y. Liu, S. Li, Y. Wang, Y. Tang, X. Li, and Y. Yang, "Soc: Semantic-assisted object cluster for referring video object segmentation," *NeurIPS*, vol. 36, 2023.
- [70] B. Miao, Z. Wu, M. Bennamoun, Y. Gao, and A. Mian, "3rd place solution for the 5th large-scale video object segmentation: Challenge—track 3: Referring video object segmentation."
- [71] W. Zhao, K. Wang, X. Chu, F. Xue, X. Wang, and Y. You, "Modeling motion with multi-modal features for text-based video segmentation," in *CVPR*, 2022, pp. 11 737–11 746.
- [72] D. Wu, X. Dong, L. Shao, and J. Shen, "Multi-level representation learning with semantic alignment for referring video object segmentation," in *CVPR*, 2022, pp. 4996–5005.
- [73] C. Liang, Y. Wu, Y. Luo, and Y. Yang, "Clawcranenet: Leveraging object-level relation for text-based video segmentation," *arXiv preprint arXiv:2103.10702*, 2021.
- [74] S. Liu, T. Hui, S. Huang, Y. Wei, B. Li, and G. Li, "Cross-modal progressive comprehension for referring segmentation," vol. 44, no. 9, pp. 4761–4775, 2021.

- [75] Z. Yang, Y. Tang, L. Bertinetto, H. Zhao, and P. H. Torr, "Hierarchical interaction network for video object segmentation from referring expressions," in *BMVC*, 2021.
- [76] Z. Ding, T. Hui, J. Huang, X. Wei, J. Han, and S. Liu, "Language-bridged spatial-temporal interaction for referring video object segmentation," in *CVPR*, 2022, pp. 4964–4973.
- [77] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu, "Video swin transformer," in *CVPR*, 2022, pp. 3202–3211.
- [78] W. Chen, D. Hong, Y. Qi, Z. Han, S. Wang, L. Qing, Q. Huang, and G. Li, "Multi-attention network for compressed video referring object segmentation," in *ACM MM*, 2022, pp. 4416–4425.
- [79] B. Yan, Y. Jiang, J. Wu, D. Wang, P. Luo, Z. Yuan, and H. Lu, "Universal instance perception as object discovery and retrieval," *arXiv preprint arXiv:2303.06674*, 2023.
- [80] X. Li, J. Wang, X. Xu, X. Li, B. Raj, and Y. Lu, "Robust referring video object segmentation with cyclic structural consensus," in *ICCV*, 2023, pp. 22 236–22 245.
- [81] D. Wu, T. Wang, Y. Zhang, X. Zhang, and J. Shen, "Onlinerefer: A simple online baseline for referring video object segmentation," in *ICCV*, 2023, pp. 2761–2770.
- [82] M. Han, Y. Wang, Z. Li, L. Yao, X. Chang, and Y. Qiao, "Htm1: Hybrid temporal-scale multimodal learning framework for referring video object segmentation," in *ICCV*, 2023, pp. 13 414–13 423.
- [83] C. Liang, Y. Wu, T. Zhou, W. Wang, Z. Yang, Y. Wei, and Y. Yang, "Rethinking cross-modal interaction from a top-down perspective for referring video object segmentation," *arXiv preprint arXiv:2106.01061*, 2021.
- [84] Y. Wang, Z. Xu, X. Wang, C. Shen, B. Cheng, H. Shen, and H. Xia, "End-to-end video instance segmentation with transformers," in *CVPR*, 2021, pp. 8741–8750.
- [85] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *ICCV*, 2021, pp. 10 012–10 022.
- [86] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [87] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, "Deformable detr: Deformable transformers for end-to-end object detection," in *ICLR*, 2021.
- [88] B. Cheng, A. Choudhuri, I. Misra, A. Kirillov, R. Girdhar, and A. G. Schwing, "Mask2former for video instance segmentation," *arXiv preprint arXiv:2112.10764*, 2021.
- [89] A. Kamath, M. Singh, Y. LeCun, G. Synnaeve, I. Misra, and N. Carion, "Mdetr-modulated detection for end-to-end multi-modal understanding," in *ICCV*, 2021, pp. 1780–1790.
- [90] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *CVPR*, 2017, pp. 2117–2125.
- [91] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "Cbam: Convolutional block attention module," in *ECCV*, 2018, pp. 3–19.
- [92] H. W. Kuhn, "The hungarian method for the assignment problem," *Nav. Res. Logist. Q.*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [93] X. Li, X. Sun, Y. Meng, J. Liang, F. Wu, and J. Li, "Dice loss for data-imbalanced nlp tasks," in *ACL*, 2020.
- [94] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *ICCV*, 2017, pp. 2980–2988.
- [95] H. Rezatofighi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, and S. Savarese, "Generalized intersection over union: A metric and a loss for bounding box regression," in *CVPR*, 2019, pp. 658–666.
- [96] L. Ye, M. Roohan, Z. Liu, and Y. Wang, "Cross-modal self-attention network for referring image segmentation," in *CVPR*, 2019, pp. 10 502–10 511.
- [97] M. Sun, J. Xiao, E. G. Lim, C. Zhao, and Y. Zhao, "Unified multi-modality video object segmentation using reinforcement learning," *IEEE Trans. Circuits Syst. Video Technol.*, 2023.
- [98] M. Gao, J. Yang, J. Han, K. Lu, F. Zheng, and G. Montana, "Decoupling multimodal transformers for referring video object segmentation," *IEEE Trans. Circuits Syst. Video Technol.*, 2023.
- [99] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 davis challenge on video object segmentation," *arXiv:1704.00675*, 2017.
- [100] J. Mao, J. Huang, A. Toshev, O. Camburu, A. L. Yuille, and K. Murphy, "Generation and comprehension of unambiguous object descriptions," in *CVPR*, 2016, pp. 11–20.
- [101] L. Yu, P. Poirson, S. Yang, A. C. Berg, and T. L. Berg, "Modeling context in referring expressions," in *ECCV*, 2016, pp. 69–85.
- [102] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.
- [103] K. Gavriluk, A. Ghodrati, Z. Li, and C. G. Snoek, "Actor and action video segmentation from a sentence," in *CVPR*, 2018, pp. 5958–5966.
- [104] H. Wang, C. Deng, J. Yan, and D. Tao, "Asymmetric cross-guided attention network for actor and action video segmentation from natural language query," in *ICCV*, 2019, pp. 3939–3948.
- [105] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *ICLR*, 2021.