

MPG-SAM 2: Adapting SAM 2 with Mask Priors and Global Context for Referring Video Object Segmentation

Fu Rong¹, Meng Lan², Qian Zhang³, Lefei Zhang^{1*}

¹National Engineering Research Center for Multimedia Software,
School of Computer Science, Wuhan University

²Hong Kong University of Science and Technology ³Horizon Robotics

Abstract

Referring video object segmentation (RVOS) aims to segment objects in a video according to textual descriptions, which requires the integration of multimodal information and temporal dynamics perception. The Segment Anything Model 2 (SAM 2) has shown great effectiveness across various video segmentation tasks. However, its application to offline RVOS is challenged by the translation of the text into effective prompts and a lack of global context awareness. In this paper, we propose a novel RVOS framework, termed MPG-SAM 2, to address these challenges. Specifically, MPG-SAM 2 employs a multimodal encoder to jointly encode video and textual features, generating semantically aligned video and text embeddings along with multimodal class tokens. A mask prior generator is devised to utilize the video embeddings and class tokens to create pseudo masks of target objects and global context. These masks are fed into the prompt encoder as dense prompts, along with multimodal class tokens as sparse prompts to generate accurate prompts for SAM 2. To provide the online SAM 2 with a global view, we propose a hierarchical global-historical aggregator, which allows SAM 2 to aggregate global and historical information of target objects at both pixel and object levels, enhancing the target representation and temporal consistency. Extensive experiments on several RVOS benchmarks demonstrate the superiority of MPG-SAM 2 and the effectiveness of the proposed modules. The code is available at <https://github.com/rongfu-dsb/MPG-SAM2>.

1. Introduction

Referring video object segmentation (RVOS) [23, 29, 32, 47] aims to segment the target objects in a video according to textual descriptions. This task, combining video segmentation [21] and language comprehension, necessitates not only proficiency in segmentation and inter-frame informa-

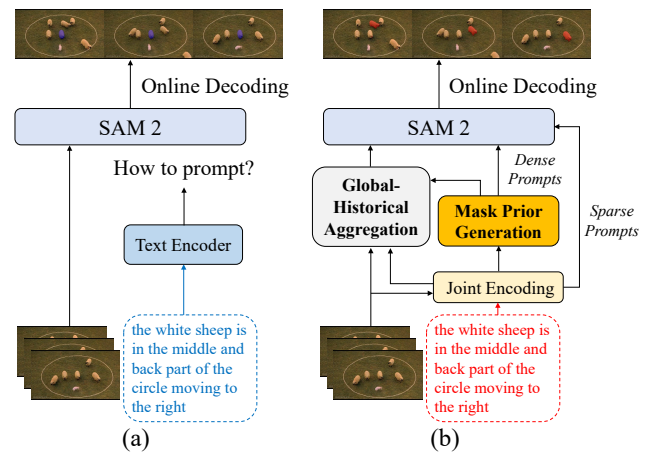


Figure 1. Comparison of two SAM 2 adaptations for RVOS. (a) vanilla SAM 2, (b) our MPG-SAM 2.

tion propagation, as in traditional video object segmentation (VOS) [1, 6, 34], but also a robust understanding of the textual reference within the broader video context. The core challenges of RVOS thus lie in efficiently aligning multimodal information and maintaining temporal consistency.

Recently, the Segment Anything Model (SAM) [19] and its variants [17, 49, 57] have demonstrated significant improvements in efficiency and accuracy for promptable segmentation in images, leveraging robust segmentation abilities and interactive prompts. SAM 2 [38] extends promptable segmentation from the image to the video domain by introducing a memory mechanism to enhance temporal consistency, achieving remarkable performance in VOS. However, its application to RVOS presents several challenges.

First, the inherent absence of textual prompts in SAM 2’s architecture hinders the delivery of accurate prompts that align with the provided textual descriptions, as shown in Fig. 1 (a). Although some efforts have explored this area, further improvement is still needed. For instance, RefSAM [25] projects text embedding into sparse and dense prompts for SAM, but the independently encoded text prompts may

*Corresponding author

not fully capture visual semantics, limiting the effective use of SAM’s segmentation capability. AL-Ref-SAM 2 [16] employs GPT-4 and Grounding DINO [27] to translate textual information into a box prompt of the target object, but the multi-stage pipeline heavily depends on the upstream model’s spatio-temporal reasoning ability, with substantial model parameters constraining deployment and inference efficiency of the model. Therefore, how to effectively align the vision-language features and provide accurate prompts to guide the decoding process is essential for adapting SAM 2 to RVOS. **Second**, the online-mode SAM 2 only has a historical view and cannot provide a global perspective for offline-mode RVOS, which may affect the global alignment of multimodal information and the temporal consistency of target objects. Consequently, effectively injecting global context information of the target objects into SAM 2 is crucial for RVOS.

To address these challenges, this paper introduces MPG-SAM 2, a novel end-to-end RVOS framework adapted from SAM 2. As illustrated in Fig. 1 (b), *our core innovation lies in generating precise prompts and global context through aligned video-text features for injection into SAM 2*. Specifically, we first employ an existing multimodal encoder to jointly encode input video and text, producing semantically aligned video and text embeddings, and multimodal class tokens. Then, we devise a novel mask prior generator, which leverages the video embeddings and multimodal class tokens to create the pseudo masks of target objects for each frame in the video, serving as dense prompts of SAM 2 that provide strong positional guidance for mask decoding. Additionally, following [58], we sent the multimodal class tokens to the prompt encoder as the sparse prompts after a multilayer perceptron (MLP). By combining the powerful dense and sparse prompts, accurate prompts are provided to the mask decoder for better performance.

To introduce the global context of target objects into SAM 2, we design a hierarchical global-historical aggregator that allows SAM 2 to aggregate global context and historical information of target objects at multiple levels before mask decoder. Here, the global context primarily comes from the global video feature generated by the mask prior generator. The aggregator consists of pixel and object-level fusion modules. In the pixel-level module, the current image feature interacts sequentially with the historical features in the memory of SAM 2 and the global context, thus enhancing the pixel-level target representation from various perspectives. Similarly, in the object-level module, the mask tokens aggregate target representation information from the global video feature and historical mask tokens in memory to generate the object tokens for the mask decoder.

Experimental results on several RVOS benchmarks demonstrate the state-of-the-art performance of our model and the effectiveness of our proposed modules. The main

contributions of this work can be summarized as follows:

- We propose a novel RVOS framework, MPG-SAM 2, adapted from SAM 2 by introducing mask prior-based dense prompt and multi-level global context fusions, achieving cutting-edge performance on several RVOS benchmarks.
- We devise a mask prior generator that leverages the global video feature and multimodal class tokens to produce pseudo masks of target objects, providing the prior position cues as dense prompts for SAM 2 to enhance the mask decoding.
- We develop a hierarchical global-historical aggregator that integrates global context and historical memory information of target objects into SAM 2 at both pixel and object levels. This module enables the online SAM 2 to have a global view and enhances the target representation and temporal consistency.

2. Related work

Referring Video Object Segmentation. RVOS attracts extensive interest as it bridges visual and linguistic domains. Early methods, such as RefVOS [2], consider RVOS as an expansion of referring image segmentation (RIS) to the video domain. URVOS [40] advances this by combining RIS and semi-supervised video object segmentation within a unified framework using attention mechanisms. In subsequent studies [12, 24, 31, 43, 45], researchers emphasize cross-modal interaction [13, 48], further enhancing RVOS performance. However, despite notable gains, the computational expense and complexity of multi-stage pipelines limit practical feasibility. In response, query-based Transformer architectures [22] offer efficient solutions with simplified yet robust frameworks. Notably, MTTR [3] and ReferFormer [47] pioneer the use of the DETR series [5, 59] in RVOS, introducing novel multimodal interaction mechanisms. Recent methods [23, 29, 32, 35, 41] refine these Transformer architectures through advanced temporal and multimodal feature integration techniques. For instance, SgMg [32] enhances ReferFormer by replacing dynamic convolution with a segmentation optimizer and uses spectral information to guide visual feature fusion. SOC [29] implements layered temporal modeling across video and object levels, achieving early multimodal fusion. More recently, Losh [55] introduces a joint prediction network for short and long sentences, reducing the over-influence of action and relational cues in segmentation. VD-IT [60] explores the potential of text-to-video diffusion models in the RVOS task, leveraging the inherent rich semantics and coherent temporal correspondences of video generation models to ensure temporal instance consistency. To tackle longer sequences and complex scenarios in the MeViS [11] dataset, DsHmp [15] proposes a static text decoupling strategy, enhancing the temporal understanding of static and dynamic

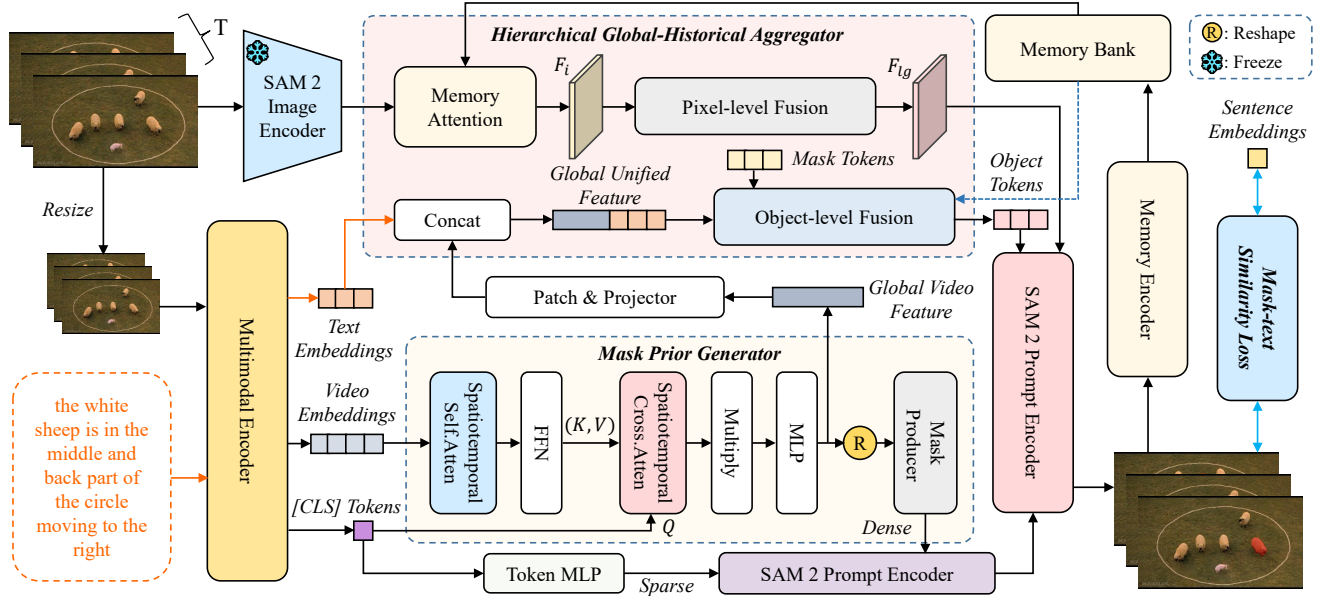


Figure 2. The overview of the proposed MPG-SAM 2. It mainly consists of four parts: the multimodal encoder, the mask prior generator, the hierarchical global-historical aggregator, and the SAM 2. The mask prior generator generates pseudo masks of target objects as the dense prompts of SAM 2 and produces the global video feature as the global context. The global-historical aggregator aggregates the target information from the global context and memory features to enhance the target representation of pixel-level image features and object-level object tokens. The mask-text similarity loss constrains the correlation between object masks and textual features.

content at both frame and object levels, thereby capturing both short-term and long-term motion insights.

Segment Anything Model. SAM [19] is an interactive segmentation model capable of generating non-semantic masks based on various prompts. Trained on a large-scale dataset, SAM exhibits strong generalization across a wide range of common objects. Several variants [17, 49, 57] have focused on improving the segmentation accuracy and computational efficiency of SAM. Moreover, SAM has found widespread application in various fields, including video tracking [8], remote sensing image interpretation [39, 42], and medical image processing [56]. Recently, SAM 2 [17] extends SAM to the video domain, achieving state-of-the-art performance. Despite its excellence in visual segmentation tasks using box, point, or mask prompts, SAM lacks language comprehension capabilities and cannot directly handle referring segmentation tasks. With the advancement of multimodal large language models (MLLMs), recent approaches [20, 37, 50] utilize MLLMs to encode textual guidance for SAM segmentation. Specifically, LISA [20] fine-tunes LLaVA [26] to generate multimodal features by extracting hidden embeddings based on specific prompts. u-LLaVA [50] extends this approach to enable joint multi-task processing at region and pixel levels. GLaMM [37] integrates the visual grounding task and incorporates language responses to provide multi-granular segmentation prompts to the model. Most recently, EVF-SAM [58] introduces

a lightweight pre-fusion architecture that employs joint visual-language encoding to generate high-quality textual prompts, leading to excellent segmentation performance.

3. Method

3.1. Overview

The overview of our proposed MPG-SAM 2 is illustrated in Fig. 2. MPG-SAM 2 consists of four primary components: the multimodal encoder, the proposed mask prior generator, the devised hierarchical global-historical aggregator, and the SAM 2. Given a video sequence $\mathcal{V} = \{I_t\}_{t=1}^T$ with T frames and its corresponding textual description $\mathcal{E} = \{e_l\}_{l=1}^L$ with L words, the multimodal encoder first performs joint encoding independently across frames, extracting the multimodal [CLS] tokens, the video patch embeddings and the text embeddings. Simultaneously, the image encoder of SAM 2 independently extracts video frame features. The mask prior generator receives the video patch embeddings and multimodal [CLS] tokens, generates the global video feature and produces prior masks for each frame. These prior masks along with the multimodal [CLS] tokens, serve as the prompts for SAM 2. The hierarchical global-historical aggregator integrates the global video feature, text embeddings, and the historical mask features and tokens from SAM 2’s memory to hierarchically enhance the target representations of pixel-level image features and

object-level object tokens. Finally, SAM 2’s decoder performs online decoding based on the provided prompts, object tokens, and current image features to obtain precise object masks for the current frame.

3.2. Feature Extraction

Aligning the visual-linguistic space is essential for RVOS. Following EVF-SAM [58] and Shared-RIS [54], we employ the unified BEiT-3 [44] encoder to jointly encode video and language features. Each input frame $I_m \in \mathbb{R}^{H_m \times W_m \times 3}$ is partitioned into non-overlapping patches $V_p \in \mathbb{R}^{N_v \times (p^2 \times 3)}$, which are subsequently projected into the feature space as $V'_p \in \mathbb{R}^{N_v \times D}$. A visual class token V_{cls} is prepended, followed by the addition of learnable positional embeddings V_{pos} , forming the final visual representation V_0 . In parallel, the textual input of length L is tokenized via XLMRobertaTokenizer [9], producing T_{tok} , with a class token T_{cls} and an end-of-sequence token T_{end} appended. Positional embeddings T_{pos} are incorporated to obtain the text representation $T_0 \in \mathbb{R}^{N_l \times D}$, where $N_l = L + 2$. The final joint representation G_0 is constructed by concatenating the visual embeddings V_0 and textual embeddings T_0 . The process is formally expressed as:

$$\begin{aligned} V_0 &= [V_{cls}, V_p] + V_{pos}, \\ T_0 &= [T_{cls}, T_{tok}, T_{end}] + T_{pos}, \\ G_0 &= [V_0; T_0] \in \mathbb{R}^{(N_v + N_l + 1) \times D}. \end{aligned} \quad (1)$$

Following multimodal fusion through multiple attention blocks, the joint visual-text embedding is fed into modality-specific feed-forward neural network (FFN) for vision and text. We ultimately obtain the joint visual-text embeddings $G \in \mathbb{R}^{T \times (N_v + N_l + 1) \times D}$ for the entire video, which are then decomposed into the multimodal [CLS] tokens $V_{cls} \in \mathbb{R}^{T \times 1 \times D}$, video patch embeddings $V \in \mathbb{R}^{T \times N_v \times D}$, and text embeddings $T \in \mathbb{R}^{T \times N_l \times D}$.

Meanwhile, the SAM 2 image encoder performs frame-independent encoding on each input video frame $I_s \in \mathbb{R}^{H_s \times W_s \times 3}$. For the i -th frame, a set of hierarchical multi-scale features is extracted for mask generation, with the image feature of final layer $F_i \in \mathbb{R}^{\frac{H_s}{16} \times \frac{W_s}{16} \times C}$ serving as the input for the subsequent decoding process, where H_s , W_s and C denote the height and width of the input image, and the channel of SAM 2 image encoder, respectively.

3.3. Mask Prior Generator

Despite the demonstrated effectiveness of semantic alignment between video patch embeddings V and text embeddings T , we identify two critical limitations in the current framework: (1) SAM 2’s video features F inherently lack linguistic context, resulting in a semantic gap with text representations, and (2) the frame-agnostic characteristic of [CLS] tokens restricts their ability to model temporal dependencies across video sequences, often leading to

segmentation artifacts including object misalignment and temporal inconsistency. To mitigate these limitations, we introduce a novel approach that generates frame-specific pseudo mask priors through the fusion of [CLS] tokens with language-enhanced video patch embeddings. These dynamically generated priors provide precise pixel-level guidance, significantly enhancing the SAM 2’s decoding process.

As illustrated in Fig. 2, we start by establishing inter-frame interaction among the video patch embeddings for each frame. Frame-agnostic video patch embeddings are temporally and spatially unfolded into video embeddings $V' \in \mathbb{R}^{(T \times N_v) \times D}$, which are subsequently fed into the multi-head self-attention layer and FFN to model the spatiotemporal context across the $T \times N_v$ dimension, learning pixel-wise inter-frame global correlations. Afterward, the [CLS] tokens for each frame are similarly flattened across the spatiotemporal dimensions, resulting in video class embeddings $V'_{cls} \in \mathbb{R}^{(T \times 1) \times D}$. To generate frame-consistent mask priors while enriching the class embeddings with visual context, we employ a multi-head cross-attention to facilitate spatiotemporal interaction between the video class embedding and the video patch embeddings. Here, the video class embeddings V'_{cls} serve as queries, and the video embeddings V' act as keys and values. The whole process could be formulated as follows:

$$\begin{aligned} V' &= \mathbf{R}(V) + \text{MHSA}(\mathbf{R}(V)), \\ V'_{cls} &= \mathbf{R}(V_{cls}) + \text{MHCA}(\mathbf{R}(V_{cls}), V') \end{aligned} \quad (2)$$

where $\mathbf{R}(\cdot)$ denotes the reshape operation, MHSA and MHCA represent multi-head self-attention and cross-attention, respectively.

The preceding steps establish inter-frame coherence in both video patch embeddings V and class embeddings V_{cls} . The class embeddings V_{cls} , jointly encoded with visual and linguistic modalities, effectively represent the intersection of visual and textual information, specifically the object denoted in the text. Therefore, after the video class embeddings V'_{cls} are broadcast to match the dimensions of the video embeddings V' , element-wise multiplication is performed. The resulting embeddings are then processed by a MLP layer to generate the global video feature V_g enriched with object information. Simultaneously, dimensionality reduction is applied to both the video and class embeddings. The video embeddings are reshaped into feature maps with dimensions $\frac{H_m}{p} \times \frac{W_m}{p}$. The class embeddings V_{cls} , representing foreground information, are multiplied element-wise with these reshaped video embeddings representing global features. This product is subsequently passed through the mask producer, which consists of a MLP layer, to generate the frame-specific mask priors $M_p \in \mathbb{R}^{T \times \frac{H_m}{p} \times \frac{W_m}{p}}$. This process could be formulated as follows:

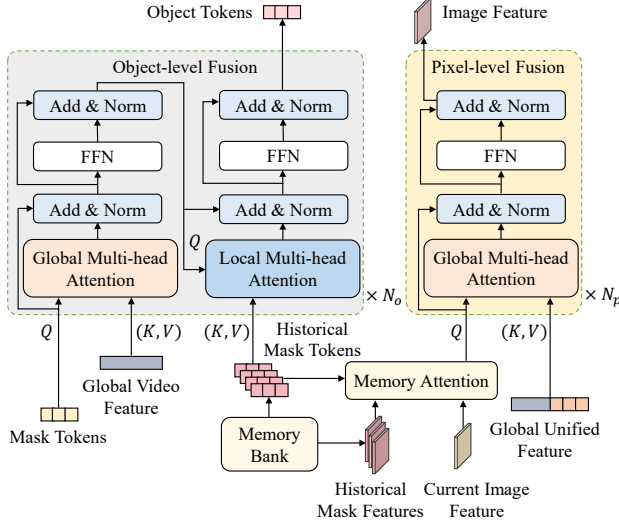


Figure 3. The structure of the hierarchical global-historical aggregator.

$$V_g = \text{MLP}(V' \cdot V'_{cls}), M_p = \text{MLP}(V_g). \quad (3)$$

Then, the global video feature V_g is input into the hierarchical global-historical aggregator to provide global object context, and the mask priors M_p are supplied as dense prompts to the prompt encoder of SAM 2.

3.4. Hierarchical Global-Historical Aggregator

Unlike VOS that primarily relies on past information, RVOS emphasizes the effective utilization of both temporal and textual cues. Inspired by [7], we propose a hierarchical global-historical aggregator that synergistically combines SAM 2’s memory mechanism with multi-level temporal modeling, enabling comprehensive integration of global context and historical segmentation results at both pixel-wise and object-centric levels. As shown in Fig. 3, the hierarchical global-historical aggregator comprises two components: the pixel-level fusion module and the object-level fusion module.

The pixel-level fusion module processes the current frame’s image feature F_i through a two-stage attention mechanism. First, memory attention is performed with historical mask features and mask tokens from the memory bank, enhancing the target representation with localized temporal information, resulting in the locally enhanced feature F_l . Subsequently, F_l undergoes global multi-head attention with the global unified features V_u , which are generated by concatenating compressed global video features V_g with corresponding text features. Although text features are encoded independently for each frame, their semantic consistency across the video sequence enables them to serve as global context. The final output of this module, the en-

hanced frame feature F_{lg} , is obtained through the add & layer norm operation and a FFN layer, effectively combining local and global information.

In the object-level fusion module, the mask tokens of the current frame T_m first engage in cross-attention with the global video feature V_g from the mask prior generator. After the add & layer norm operation and a FFN layer, the globally enhanced mask tokens T_{mg} undergo another cross-attention operation with historical mask tokens from the memory bank. This process yields the final object tokens T_{mgl} that incorporate both global context and historical information through another round of add & layer norm operation and a FFN layer.

The enhanced object tokens T_{mgl} are then fed into the mask decoder, replacing the original mask tokens, while the enhanced frame features F_{lg} are simultaneously processed. The object queries, composed of sparse prompts and enhanced object tokens, along with the text-enriched frame features F_{lg} , collectively enable a more accurate mask decoding process. This dual enhancement strategy ensures that both the object queries and video features are infused with comprehensive temporal and textual information, significantly improving the segmentation accuracy.

3.5. SAM 2 Prompt Encoder and Mask Decoder

In MPG-SAM 2, the prompt encoder processes two distinct types of prompts: sparse prompts and dense prompts. For the sparse prompts, following EVF-SAM [58], we project the [CLS] tokens V_{cls} from the multimodal encoder through a token MLP and then concatenate them with zero-initialized sparse tokens to form the sparse prompts. The dense prompts are generated by upsampling the mask priors M_p from the mask prior generator through linear interpolation, ensuring spatial alignment with SAM 2’s feature dimensions.

The SAM 2 mask decoder is designed to utilize both prompt types synergistically: (1) the sparse prompts, concatenated with object tokens, form the object-level queries for the decoder, while (2) the dense pixel-level prompts are element-wise added to the video frame features, providing direct guidance from high-resolution feature layers.

While the original SAM 2 architecture does not inherently support concurrent processing of both prompt types, our modifications enable effective integration of multimodal sparse embeddings and dense pseudo mask priors. This architectural enhancement significantly improves the mask decoder’s capability to interpret text-referenced objects, which is crucial for RVOS performance.

3.6. Training Loss

MPG-SAM 2 employs an overall loss function similar to that of [47] to constrain the predicted mask as follows:

$$\mathcal{L} = \lambda_{dice} \mathcal{L}_{dice} + \lambda_{focal} \mathcal{L}_{focal} + \lambda_{sim} \mathcal{L}_{sim}, \quad (4)$$

here, \mathcal{L}_{dice} is the DICE loss [33], \mathcal{L}_{focal} represents the binary mask focal loss and \mathcal{L}_{sim} denotes the mask-text similarity loss, detailed as follows.

Mask-text Similarity Loss. In the context of the RVOS task, evaluating the similarity between the segmentation mask and the ground truth mask is essential. However, in addition to the traditional mask-based evaluation criteria, a mask-text loss function can also be introduced to assess the segmentation results. The similarity between the text and mask can serve as an additional evaluation metric. Specifically, We use sentence embeddings T_s as the abstract representation of the text embeddings T , which is dimensionally compressed to a singular scalar through MLP layers, subsequently expanded to match the dimensions of the mask, and utilized as an output to calculate the pixel-level similarity S_{tp} between the text and the predicted mask M_{pre} , as well as the pixel-level similarity S_{tg} between the text and the ground truth mask M_{gt} :

$$\begin{aligned} S_{tp} &= f_{cos}(\text{MLP}(T_s), M_{pre}), \\ S_{tg} &= f_{cos}(\text{MLP}(T_s), M_{gt}), \end{aligned} \quad (5)$$

where f_{cos} denotes the cosine similarity function. Subsequently, MSE loss is employed to enforce pixel-level constraints between S_{tp} and S_{tg} :

$$\mathcal{L}_{sim} = \frac{1}{N} \sum_{i=1}^N (S_{tp_i} - S_{tg_i})^2, \quad (6)$$

where N represents the number of pixels in the whole video.

4. Experiments

4.1. Datasets and Metrics

Datasets. The experiments are performed on several key RVOS datasets: Ref-YouTube-VOS [40], MeViS [11] and Ref-DAVIS17 [18]. Ref-YouTube-VOS is a widely recognized and large-scale dataset in the field of RVOS, containing 3471 videos with 12913 expressions in the training set and 202 videos with 2096 expressions in the validation set. MeViS, a newly established dataset, focuses on motion analysis, consisting of 2,006 videos with 28,570 annotations. Ref-DAVIS17 builds on the DAVIS17 [36] dataset with additional linguistic annotations for diverse objects, comprising 90 videos.

Evaluation Metrics. We adhere to the standard evaluation framework outlined in [40], employing metrics such as region similarity \mathcal{J} , contour accuracy \mathcal{F} , and their combined average $\mathcal{J}\&\mathcal{F}$ to evaluate our model on the validation sets of Ref-YouTube-VOS, MeViS and Ref-DAVIS17. Due to the absence of publicly accessible ground truth annotations for the Ref-YouTube-VOS and MeViS validation set, we utilize the official server to submit our predictions and obtain the evaluation results.

4.2. Implementation Details

Model Settings. We initialize the relevant modules of SAM 2 and the multimodal encoder using the SAM 2-Hiera-Large [38] and BEiT-3-Large [44] pre-trained weights. For feature parsing, each image is resized to resolutions of 1024×1024 and 224×224 , serving as the input to SAM 2 image encoder with an output dimension C of 256, and the multimodal encoder with an output dimension D of 1024, respectively. In the hierarchical global-historical aggregator, the patch size p_g for the global video feature in pixel-level fusion is set to 2. Both the pixel-level fusion layer number N_p and object-level fusion layer number N_o are set to 1. The memory bank is configured similarly to SAM 2 [38], with a maximum storage capacity of 7 historical mask features and 16 mask tokens.

Training Details. Experiments are conducted on 8 NVIDIA A800 GPUs for the MeViS [11] dataset due to its high memory requirements stemming from dataset-specific configuration and on 8 NVIDIA GeForce RTX 4090 GPUs for the remaining datasets. The MeViS dataset experiments follow settings similar to [15], with 8 frames as input for training, and are trained directly on this dataset without any pre-training on RefCOCO+/g [30, 53]. Training is performed for 6 epochs using the AdamW optimizer [28], set at a learning rate of $2e-6$. For the Ref-YouTube-VOS [40] and Ref-DAVIS17 [18] datasets, we adopt an approach similar to [32, 47], first pre-training on the RefCOCO+/g [30, 53] for 10 epochs, followed by fine-tuning on Ref-YouTube-VOS for 6 epochs. The batch size during pre-training is set to 8, with a learning rate of $1e-5$, while fine-tuning use a batch size of 1, a learning rate of $2e-6$, and 5 frames as input. To better align with the pre-trained parameters of other modules, we set a learning rate of $5e-5$ for both the hierarchical global-historical aggregator and the mask prior generator. The trained model is then validated on the Ref-DAVIS17 dataset without additional training. The loss weights for different losses are set as follows: $\lambda_{focal} = 2$, $\lambda_{dice} = 5$, and $\lambda_{sim} = 2$.

4.3. Comparison with State-of-the-Art Methods

Ref-YouTube-VOS & Ref-DAVIS17 sets. We compare our MPG-SAM 2 approach with several state-of-the-art methods, as shown in Tab. 1. Our approach outperforms all existing methods on both datasets. On the Ref-YouTube-VOS [40] dataset, we achieve 73.9% $\mathcal{J}\&\mathcal{F}$, surpassing the best methods DsHmp [15] by 6.8% $\mathcal{J}\&\mathcal{F}$ and LoSh [55] by 6.7% $\mathcal{J}\&\mathcal{F}$. Even compared to methods using additional training data, our model performs competitively, exceeding MUTR [52] by 5.5% $\mathcal{J}\&\mathcal{F}$. On the Ref-DAVIS [18] dataset, our method achieves 72.4% $\mathcal{J}\&\mathcal{F}$, outperforming VD-IT [60] by 3.0% $\mathcal{J}\&\mathcal{F}$ and surpassing VISA [51] by 2.0% $\mathcal{J}\&\mathcal{F}$, which is trained on additional datasets. All comparison methods use the optimal configuration to high-

Table 1. Comparison with state-of-the-art methods on the Ref-YouTube-VOS and Ref-DAVIS17 datasets. The best results are highlighted in bold, and the second best results are underlined.

Method	Reference	Ref-YouTube-VOS			Ref-DAVIS17		
		$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
ReferFormer [47]	CVPR'22	62.9	61.3	64.6	61.1	58.1	64.1
OnlineRefer [46]	ICCV'23	62.9	61.0	64.7	62.4	59.1	65.6
HTML [14]	ICCV'23	63.4	61.5	65.2	62.1	59.2	65.1
SgMg [32]	ICCV'23	65.7	63.9	67.4	63.3	60.6	66.0
TempCD [41]	ICCV'23	65.8	63.6	68.0	64.6	61.6	67.6
SOC [29]	NIPS'23	66.0	64.1	67.9	64.2	61.0	67.4
LoSh [55]	CVPR'24	67.2	65.4	69.0	64.3	61.8	66.8
DsHmp [15]	CVPR'24	67.1	65.0	69.1	64.9	61.7	68.1
MUTR [52]	AAAI'24	<u>68.4</u>	<u>66.4</u>	70.4	68.0	64.8	71.3
VD-IT [60]	ECCV'24	66.5	64.4	68.5	69.4	66.2	72.6
VISA [51]	ECCV'24	63.0	61.4	64.7	<u>70.4</u>	<u>67.0</u>	<u>73.8</u>
MPG-SAM 2	-	73.9	71.7	76.1	72.4	68.8	76.0

Table 2. Comparison with state-of-the-art methods on the MeViS dataset. Our model achieves the best performance.

Method	Reference	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
URVOS [40]	ECCV'20	27.8	25.7	29.9
LBDT [12]	CVPR'22	29.3	27.8	30.8
MTTR [4]	CVPR'22	30.0	28.8	31.2
ReferFormer [47]	CVPR'22	31.0	29.8	32.2
VLT+TC [10]	TPAMI'22	35.5	33.6	37.3
LMPM [11]	ICCV'23	37.2	34.2	40.2
VISA [51]	ECCV'24	44.5	41.8	47.1
DsHmp [15]	CVPR'24	46.4	43.0	49.8
MPG-SAM 2	-	53.7	50.7	56.7

light model performance.

MeViS set. We also conduct comparative experiments between our MPG-SAM 2 and existing methods including URVOS [40], LBDT [12], MTTR [4], ReferFormer [47], VLT+TC [10], LMPM [11], VISA [51] and DsHmp [15], on the MeViS [11] dataset, with results documented in Tab. 2. On this dataset, our method achieves a $\mathcal{J}\&\mathcal{F}$ score of 53.7%, surpassing the current state-of-the-art method DsHmp by 7.3% $\mathcal{J}\&\mathcal{F}$, demonstrating the effectiveness of our approach in leveraging temporal information.

Fig. 4 presents the visual comparison of our MPG-SAM 2 model with SgMg [32] on Ref-YouTube-VOS dataset. The results clearly demonstrate that MPG-SAM 2 consistently surpasses SgMg, especially in prediction accuracy and maintaining frame-to-frame consistency.

Although MPG enables spatiotemporal self-interaction on video embeddings, its computational and memory overhead remains acceptable due to BEiT-3's small embedding size (14×14). MPG and HGA together add 1.3G to memory and introduce 21M parameters, which is reasonable. For a

Table 3. Ablation study of different components of MPG-SAM 2 on Ref-YouTube-VOS dataset.

Method	\mathcal{L}_{sim}	MPG	HGA	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
Baseline				69.4	67.5	71.3
MPG-SAM 2	✓			70.3	68.3	72.2
MPG-SAM 2	✓	✓		71.9	69.8	73.9
MPG-SAM 2	✓		✓	72.3	70.1	74.4
MPG-SAM 2	✓	✓	✓	73.9	71.7	76.1

detailed analysis of the overall model parameters, please refer to the appendix.

4.4. Model Analysis

In this section, we conduct comprehensive ablation experiments to examine the effects of the key components of our MPG-SAM 2 and the influence of various model configurations. All experiments are carried out using the Ref-YouTube-VOS dataset.

Components Analysis. To explore the influence of the key components of our model, we first construct a **baseline** model consisting solely of SAM 2 and the multimodal encoder. We improve the video baseline over image-based referring segmentation by (1) providing per-frame prompts (P_{per}) instead of only the first frame and (2) using the SAM 2 memory mechanism (Mem) instead of frame-independent segmentation. As shown in Tab. 4, these enhancements lead to the baseline by 10.4% $\mathcal{J}\&\mathcal{F}$ and 3.2% $\mathcal{J}\&\mathcal{F}$, respectively.

After that, as depicted in Tab. 3, we first add the mask-text similarity loss (\mathcal{L}_{sim}) to the baseline, resulting in a 0.9% $\mathcal{J}\&\mathcal{F}$ improvement. Building on this, we introduce the mask prior generator (MPG), and the special MPG-SAM 2 achieves 71.9% $\mathcal{J}\&\mathcal{F}$, which is 1.6% higher than the previous model. While only the hierarchical global-historical aggregator (HGA) is imposed on the baseline with \mathcal{L}_{sim} , the $\mathcal{J}\&\mathcal{F}$ score of the special MPG-SAM 2 reaches 72.3%, indicating a 2.0% improvement and further validating the module's superiority. With all components integrated, our MPG-SAM 2 delivers the best performance of 73.9% $\mathcal{J}\&\mathcal{F}$. Note that the memory mechanism is excluded from HGA during performance calculations; when the MPG module is omitted, only the mask prior is not generated, while the global video feature production remains.

Mask Prior Generator. In this section, we investigate different spatiotemporal interaction forms within the mask prior generator (MPG), with results presented in Tab. 4. When the overall spatiotemporal self-interaction (S_{si}) of video embeddings is omitted, MPG-SAM 2 experiences a performance drop of 0.6% $\mathcal{J}\&\mathcal{F}$. Similarly, when the spatiotemporal cross-modal interaction (S_{ci}) between the [CLS] tokens and video embeddings is excluded, the model




Figure 4. Visualization result on Ref-YouTube-VOS. (a) SgMg [32], (b) MPG-SAM 2. MPG-SAM 2 effectively ensures the global spatiotemporal consistency of masks, reducing objects misalignment and drift.

Table 4. Model analysis of different settings in MPG-SAM 2.

Method	Settings	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
Settings of Baseline				
Baseline	w/o P_{per}	59.0	57.5	60.6
Baseline	w/o Mem	66.2	64.5	67.9
Baseline	w P_{per} & Mem	69.4	67.5	71.3
Interaction Mode of MPG				
MPG-SAM 2	w/o S_{si}	73.3	71.2	75.4
MPG-SAM 2	w/o S_{ci}	73.5	71.4	75.6
MPG-SAM 2	w S_{si} & S_{ci}	73.9	71.7	76.1
Patch Size of HGA				
MPG-SAM 2	1	73.2	71.0	75.4
MPG-SAM 2	2	73.9	71.7	76.1
MPG-SAM 2	4	73.6	71.4	75.7
Global Video Feature of HGA				
MPG-SAM 2	Vanilla	73.1	71.1	75.1
MPG-SAM 2	Masked	73.9	71.7	76.1

performance decreases by 0.4% $\mathcal{J}\&\mathcal{F}$. These findings highlight the importance of understanding video information from a global spatiotemporal perspective during the mask prompt generation process.

Hierarchical Global-Historical Aggregator. The impact of different settings of hierarchical global-historical aggregator (HGA) is also worth noting. First, we examine the effect of patch size p_g in the pixel-level fusion process for the global video feature, experimenting with patch sizes 1, 2, and 4. As shown in Tab. 4, the patch size of 2 yields the best model performance. This configuration achieves an optimal balance between critical and redundant information during pixel-level fusion, allowing for more effective integration of global context.

Additionally, we explore different configurations for the global video feature input to the hierarchical global-historical aggregator. One straightforward approach named “Vanilla” employs the video embeddings directly output by the multimodal encoder as the global feature. An alternative approach dubbed “Masked” leverages a video feature enriched with mask prior information. Results in Tab. 4 indicate that video features containing mask information are more beneficial for global context integration, as they better emphasize global frame mask information to guide the segmentation of the current frame.

5. Conclusion

In this paper, we present MPG-SAM 2, an innovative end-to-end framework for RVOS, to address the challenges of adapting SAM 2 to the RVOS task. Our approach utilizes a unified multimodal encoder to jointly encode video and textual features, generating semantically aligned video and text embeddings, along with multimodal class tokens. The video embeddings and class tokens are employed by a mask prior generator to create pseudo masks of target objects, offering strong positional cues as dense prompts for SAM 2’s mask decoder. To address SAM 2’s lack of global context awareness in offline RVOS, we introduce a hierarchical global-historical aggregator. This enables SAM 2 to integrate global context and historical information of target objects at both pixel and object levels, enhancing the target representation and temporal consistency. Extensive experiments on several RVOS benchmarks demonstrate the superiority of our MPG-SAM 2 over state-of-the-art methods and validate the effectiveness of our proposed modules.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62431020, in part by the Foundation for Innovative Research Groups of Hubei Province under Grant 2024AFA017, in part by the Fundamental Research Funds for the Central Universities under Grant 2042025kf0030.

References

- [1] Maksym Bekuzarov, Ariana Bermudez, Joon-Young Lee, and Hao Li. Xmem++: Production-level video segmentation from few annotated frames. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 635–644, 2023. 1
- [2] Miriam Bellver, Carles Ventura, Carina Silberer, Ioannis Kazakos, Jordi Torres, and Xavier Giro-i Nieto. Refvos: A closer look at referring expressions for video object segmentation. *arXiv preprint arXiv:2010.00263*, 2020. 2
- [3] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multi-modal transformers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4985–4995, 2022. 2
- [4] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multi-modal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4985–4995, 2022. 7
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 213–229, 2020. 2
- [6] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European Conference on Computer Vision*, pages 640–658. Springer, 2022. 1
- [7] Ho Kei Cheng, Seoung Wug Oh, Brian Price, Joon-Young Lee, and Alexander Schwing. Putting the object back into video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3151–3161, 2024. 5
- [8] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023. 3
- [9] A Conneau. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*, 2019. 4
- [10] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16321–16330, 2021. 7
- [11] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, and Chen Change Loy. Mevis: A large-scale benchmark for video segmentation with motion expressions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2694–2703, 2023. 2, 6, 7
- [12] Zihan Ding, Tianrui Hui, Junshi Huang, Xiaoming Wei, Jizhong Han, and Si Liu. Language-bridged spatial-temporal interaction for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4964–4973, 2022. 2, 7
- [13] Hao Fang, Peng Wu, Yawei Li, Xinxin Zhang, and Xiankai Lu. Unified embedding alignment for open-vocabulary video instance segmentation. In *European Conference on Computer Vision*, pages 225–241. Springer, 2024. 2
- [14] Mingfei Han, Yali Wang, Zhihui Li, Lina Yao, Xiaojun Chang, and Yu Qiao. Htm: Hybrid temporal-scale multimodal learning framework for referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13414–13423, 2023. 7
- [15] Shuting He and Henghui Ding. Decoupling static and hierarchical motion perception for referring video segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13332–13341, 2024. 2, 6, 7
- [16] Shaofei Huang, Rui Ling, Hongyu Li, Tianrui Hui, Zongheng Tang, Xiaoming Wei, Jizhong Han, and Si Liu. Unleashing the temporal-spatial reasoning capacity of gpt for training-free audio and language referenced video object segmentation. *arXiv preprint arXiv:2408.15876*, 2024. 2
- [17] Lei Ke, Mingqiao Ye, Martin Danelljan, Yu-Wing Tai, Chi-Keung Tang, Fisher Yu, et al. Segment anything in high quality. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3
- [18] Anna Khoreva, Anna Rohrbach, and Bernt Schiele. Video object segmentation with language referring expressions. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 123–141, 2018. 6
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 1, 3
- [20] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9579–9589, 2024. 3
- [21] Meng Lan, Jing Zhang, Lefei Zhang, and Dacheng Tao. Learning to learn better for video object segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1205–1212, 2023. 1
- [22] Meng Lan, Fu Rong, Hongzan Jiao, Zhi Gao, and Lefei Zhang. Language query-based transformer with multiscale cross-modal alignment for visual grounding on remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–13, 2024. 2
- [23] Meng Lan, Fu Rong, Zuchao Li, Wei Yu, and Lefei Zhang. Bidirectional correlation-driven inter-frame inter-

- action transformer for referring video object segmentation. *Pattern Recognition*, 153:110535, 2024. 1, 2
- [24] Dezhuang Li, Ruoqi Li, Lijun Wang, Yifan Wang, Jinqing Qi, Lu Zhang, Ting Liu, Qingquan Xu, and Huchuan Lu. You only infer once: Cross-modal meta-transfer for referring video object segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1297–1305, 2022. 2
- [25] Yonglin Li, Jing Zhang, Xiao Teng, Long Lan, and Xinwang Liu. Refsam: Efficiently adapting segmenting anything model for referring video object segmentation. *arXiv preprint arXiv:2307.00997*, 2024. 1
- [26] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 3
- [27] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 2
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 6
- [29] Zhuoyan Luo, Yicheng Xiao, Yong Liu, Shuyan Li, Yitong Wang, Yansong Tang, Xiu Li, and Yujiu Yang. Soc: Semantic-assisted object cluster for referring video object segmentation. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2, 7
- [30] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016. 6
- [31] Bruce McIntosh, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. Visual-textual capsule routing for text-based video segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9942–9951, 2020. 2
- [32] Bo Miao, Mohammed Bennamoun, Yongsheng Gao, and Ajmal Mian. Spectrum-guided multi-granularity referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 920–930, 2023. 1, 2, 6, 7, 8
- [33] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 565–571, 2016. 6
- [34] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9226–9235, 2019. 1
- [35] Feiyu Pan, Hao Fang, Fangkai Li, Yanyu Xu, Yawei Li, Luca Benini, and Xiankai Lu. Semantic and sequential alignment for referring video object segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19067–19076, 2025. 2
- [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. 6
- [37] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13009–13018, 2024. 3
- [38] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 1, 6
- [39] Fu Rong, Meng Lan, Qian Zhang, and Lefei Zhang. Customized sam 2 for referring remote sensing image segmentation. *arXiv preprint arXiv:2503.07266*, 2025. 3
- [40] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 208–223, 2020. 2, 6, 7
- [41] Jiajin Tang, Ge Zheng, and Sibeil Yang. Temporal collection and distribution for referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15466–15476, 2023. 2, 7
- [42] Di Wang, Jing Zhang, Bo Du, Minqiang Xu, Lin Liu, Dacheng Tao, and Liangpei Zhang. Samrs: Scaling-up remote sensing segmentation dataset with segment anything model. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [43] Hao Wang, Cheng Deng, Junchi Yan, and Dacheng Tao. Asymmetric cross-guided attention network for actor and action video segmentation from natural language query. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3939–3948, 2019. 2
- [44] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19175–19186, 2023. 4, 6
- [45] Dongming Wu, Xingping Dong, Ling Shao, and Jianbing Shen. Multi-level representation learning with semantic alignment for referring video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4996–5005, 2022. 2
- [46] Dongming Wu, Tiancai Wang, Yuang Zhang, Xiangyu Zhang, and Jianbing Shen. Onlinerefer: A simple online baseline for referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2761–2770, 2023. 7
- [47] Jiannan Wu, Yi Jiang, Peize Sun, Zehuan Yuan, and Ping Luo. Language as queries for referring video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4974–4984, 2022. 1, 2, 5, 6, 7

- [48] Peng Wu, Xiankai Lu, Hao Hu, Yongqin Xian, Jianbing Shen, and Wenguan Wang. Logiczsl: Exploring logic-induced representation for compositional zero-shot learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 30301–30311, 2025. [2](#)
- [49] Yunyang Xiong, Bala Varadarajan, Lemeng Wu, Xiaoyu Xiang, Fanyi Xiao, Chenchen Zhu, Xiaoliang Dai, Dilin Wang, Fei Sun, Forrest Iandola, et al. Efficientsam: Leveraged masked image pretraining for efficient segment anything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16111–16121, 2024. [1](#), [3](#)
- [50] Jinjin Xu, Liwu Xu, Yuzhe Yang, Xiang Li, Fanyi Wang, Yanchun Xie, Yi-Jie Huang, and Yaqian Li. u-llava: Unifying multi-modal tasks via large language model. *arXiv preprint arXiv:2311.05348*, 2023. [3](#)
- [51] Cilin Yan, Haochen Wang, Shilin Yan, Xiaolong Jiang, Yao Hu, Guoliang Kang, Weidi Xie, and Efstratios Gavves. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*, pages 98–115. Springer, 2024. [6](#), [7](#)
- [52] Shilin Yan, Renrui Zhang, Ziyu Guo, Wenchao Chen, Wei Zhang, Hongyang Li, Yu Qiao, Hao Dong, Zhongjiang He, and Peng Gao. Referred by multi-modality: A unified temporal transformer for video object segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6449–6457, 2024. [6](#), [7](#)
- [53] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 69–85. Springer, 2016. [6](#)
- [54] Seonghoon Yu, Ilchae Jung, Byeongju Han, Taeoh Kim, Yunho Kim, Dongyoon Wee, and Jeany Son. A simple baseline with single-encoder for referring image segmentation. *arXiv preprint arXiv:2408.15521*, 2024. [4](#)
- [55] Linfeng Yuan, Miaoqing Shi, Zijie Yue, and Qijun Chen. Losh: Long-short text joint prediction network for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14001–14010, 2024. [2](#), [6](#), [7](#)
- [56] Wenxi Yue, Jing Zhang, Kun Hu, Yong Xia, Jiebo Luo, and Zhiyong Wang. Surgicalsam: Efficient class promptable surgical instrument segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 6890–6898, 2024. [3](#)
- [57] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. [1](#), [3](#)
- [58] Yuxuan Zhang, Tianheng Cheng, Rui Hu, Heng Liu, Longjin Ran, Xiaoxin Chen, Wenyu Liu, Xinggang Wang, et al. Evsam: Early vision-language fusion for text-prompted segment anything model. *arXiv preprint arXiv:2406.20076*, 2024. [2](#), [3](#), [4](#), [5](#)
- [59] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020. [2](#)
- [60] Zixin Zhu, Xuelu Feng, Dongdong Chen, Junsong Yuan, Chunming Qiao, and Gang Hua. Exploring pre-trained text-to-video diffusion models for referring video object segmentation. In *European Conference on Computer Vision*, pages 452–469. Springer, 2024. [2](#), [6](#), [7](#)