

# Referred by Multi-Modality: A Unified Temporal Transformer for Video Object Segmentation

Shilin Yan<sup>1,2\*</sup>, Renrui Zhang<sup>2,3\*</sup>, Ziyu Guo<sup>3\*</sup>, Wenchao Chen<sup>1</sup>, Wei Zhang<sup>1†</sup>, Hongyang Li<sup>2†</sup>  
Yu Qiao<sup>2</sup>, Hao Dong<sup>4,5</sup>, Zhongjiang He<sup>6</sup>, Peng Gao<sup>2†</sup>

<sup>1</sup>School of Computer Science, Fudan University

<sup>2</sup>Shanghai Artificial Intelligence Laboratory

<sup>3</sup>The Chinese University of Hong Kong

<sup>4</sup>School of CS, Peking University

<sup>5</sup>PKU-agibot Lab

<sup>6</sup>China Telecom Corporation Ltd. Data&AI Technology Company

yansl21@m.fudan.edu.cn, {zhangrenrui, guoziyu, gaopeng}@pjlab.org.cn, weizh@fudan.edu.cn

## Abstract

Recently, video object segmentation (VOS) referred by multi-modal signals, e.g., language and audio, has evoked increasing attention in both industry and academia. It is challenging for exploring the semantic alignment within modalities and the visual correspondence across frames. However, existing methods adopt separate network architectures for different modalities, and neglect the inter-frame temporal interaction with references. In this paper, we propose **MUTR**, a Multi-modal Unified Temporal transformer for Referring video object segmentation. With a unified framework for the first time, MUTR adopts a DETR-style transformer and is capable of segmenting video objects designated by either text or audio reference. Specifically, we introduce two strategies to fully explore the temporal relations between videos and multi-modal signals. Firstly, for low-level temporal aggregation before the transformer, we enable the multi-modal references to capture multi-scale visual cues from consecutive video frames. This effectively endows the text or audio signals with temporal knowledge and boosts the semantic alignment between modalities. Secondly, for high-level temporal interaction after the transformer, we conduct inter-frame feature communication for different object embeddings, contributing to better object-wise correspondence for tracking along the video. On Ref-YouTube-VOS and AVSBench datasets with respective text and audio references, MUTR achieves **+4.2%** and **+8.7%**  $\mathcal{J}$ & $\mathcal{F}$  improvements to *state-of-the-art* methods, demonstrating our significance for unified multi-modal VOS. Code is released at <https://github.com/OpenGVLab/MUTR>.

## Introduction

Multi-modal video object segmentation (VOS) aims to track and segment particular object instances across the video sequence referred by a given multi-modal signal, including referring video object segmentation (RVOS) with language reference, and audio-visual video object segmentation (AV-VOS) with audio reference. Different from the vanilla

VOS (Xu et al. 2018; Yan et al. 2023) with only visual information, the multi-modal VOS is more challenging and in urgent demand, which requires a comprehensive understanding of different modalities and their temporal correspondence across frames.

There exist two main challenges in multi-modal VOS. Firstly, it requires to not only explore the rich spatial-temporal consistency in a video, but also align the multi-modal semantics among image, language, and audio. Current approaches mainly focus on the visual-language or visual-audio modal fusion within independent frames, simply by cross-modal attention (Chen et al. 2019; Hu et al. 2020; Shi et al. 2018) or dynamic convolutions (Margffoy-Tuay et al. 2018) for feature interaction. This, however, neglects the multi-modal temporal information across frames, which is significant for consistent object segmentation and tracking along the video. Secondly, for the given references of two modalities, language and audio, existing works adopt different architecture designs and training strategies to separately tackle their modal-specific characteristics. Therefore, a powerful and unified framework for multi-modal VOS still remains an open question.

To address these challenges, we propose **MUTR**, a Multi-modal Unified Temporal transformer for Referring video object segmentation. Our approach, for the first time, presents a generic framework for both language and audio references, and enhances the interaction between temporal frames and multi-modal signals. In detail, we adopt a DETR-like (Carion et al. 2020) encoder-decoder transformer, which serves as the basic architecture to process visual information within different frames. On top of this, we introduce two attention-based modules respectively for low-level multi-modal temporal aggregation (MTA), and high-level multi-object temporal interaction (MTI). Firstly before the transformer, we utilize the encoded multi-modal references as queries to aggregate informative visual and temporal features via the MTA module. We concatenate the visual features of adjacent frames and adopt sequential attention blocks for multi-modal tokens to progressively capture temporal visual cues of different image scales. This contributes to better low-level cross-modal alignment

\*These authors contributed equally.

†Corresponding Author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

and temporal consistency. Then, we regard the multi-modal tokens after MTA as object queries and feed them into the transformer for frame-wise decoding. After that, we apply the MTI module to conduct inter-frame object-wise interaction, and maintain a set of video-wise query representations for associating objects across frames inspired by (Heo et al. 2022). Such a module enhances the instance-level temporal communication and benefits the visual correspondence for segmenting the same object in a video. Finally, we utilize a segmentation head following previous works (Wu et al. 2022, 2021) to output the final object mask referred by multi-modality input.

To evaluate our effectiveness, we conduct extensive experiments on several popular benchmarks for multi-modal VOS. RVOS with language reference (Ref-YouTube-VOS (Seo, Lee, and Han 2020) and Ref-DAVIS 2017 (Khoreva, Rohrbach, and Schiele 2019)), and one benchmark for AV-VOS with audio reference (AVSBench (Zhou et al. 2022)). On Ref-YouTube-VOS (Seo, Lee, and Han 2020) and Ref-DAVIS 2017 (Khoreva, Rohrbach, and Schiele 2019) with language references, MUTR surpasses the state-of-the-art method ReferFormer (Wu et al. 2022) by +4.2% and +4.1%  $\mathcal{J}\&\mathcal{F}$  scores, respectively. On AV-VOS (Zhou et al. 2022) with audio references, we also outperform Baseline (Zhou et al. 2022) by +8.7%  $\mathcal{J}\&\mathcal{F}$  score.

Overall, our contributions are summarized as follows:

- For the first time, we present a unified transformer architecture, MUTR, to tackle video object segmentation referred by multi-modal inputs, i.e., language and audio.
- To better align the temporal information with multi-modal signals, we propose two attention-based modules, MTA and MTI, respectively for low-level multi-scale aggregation and high-level multi-object interaction, achieving superior cross-modal understanding in a video.
- On benchmarks of two modalities, our approach both achieves state-of-the-art results, e.g., +4.2 % and +4.1%  $\mathcal{J}\&\mathcal{F}$  for Ref-YouTube-VOS and Ref-DAVIS 2017, +8.7%  $\mathcal{J}\&\mathcal{F}$  for AV-VOS. This fully indicates the significance and generalization ability of MUTR.

## Related Work

**Referring video object segmentation (R-VOS).** R-VOS introduces the language expression for target object tracking and segmentation, following the trend of vision-language learning (Zhang et al. 2022, 2023b; Zhu et al. 2023; Fang et al. 2023). Existing R-VOS methods can be broadly classified into three categories. One of the most straightforward ideas is to apply referring image segmentation methods (Ding et al. 2021; Yang et al. 2022; Wang et al. 2022) independently to video frames, such as RefVOS (Bellver et al. 2020). Obviously, it disregards the temporal information, which makes it difficult to process common video challenges like object disappearance in reproduction. Another approach involves propagating the target mask detected from key frame and selecting the object to be segmented based on a visual grounding model (Kamath et al. 2021; Luo et al. 2020). Although it applies the temporal information to some extent, its complex multi-stage training approach is not desirable. The recent

work MTTR (Botach, Zheltonozhskii, and Baskin 2022) and ReferFormer (Wu et al. 2022) have employed query-based mechanisms. Nevertheless, they are end-to-end frameworks, they perform R-VOS task utilizing image-level segmentation. Contrastly, our unified framework fully explores video-level visual-attended language information for low-level temporal aggregation.

**Audio-visual video object segmentation (AV-VOS).** Inspired by recent multi-modality efforts (Zhang et al. 2023a; Gao et al. 2023; Lin et al. 2023; Wang et al. 2023; Guo et al. 2023; Han et al. 2023b,a), AV-VOS is proposed for predicting pixel-level individual positions based on a given sound signal. There is little previous work on audio-visual video object segmentation. Until recently (Zhou et al. 2022) proposed the audio-visual video object segmentation dataset. Different from it, (Mo and Tian 2023) is based on the recent visual foundation model Segment Anything Model (Kirillov et al. 2023; Zhang et al. 2023c) to achieve audio-visual segmentation. However, all of them lack the temporal alignment between multi-modal information.

## Method

In this section, we illustrate the details of our MUTR for multi-modal video object segmentation. We first describe the overall pipeline in Section . Then, in Section and Section , we respectively elaborate on the proposed designs of the multi-scale temporal aggregation module (MTA), and multi-object temporal interaction module (MTI).

### Overall Pipeline

The overall pipeline of MUTR is shown in Figure 1. We adopt a DETR-based (Carion et al. 2020) transformer as our basic architecture, including a visual backbone, a visual encoder and a decoder, on top of which, two modules MTA and MTI are proposed for temporal multi-modal interaction. In this section, we successively introduce the pipeline of MUTR for video object segmentation.

**Feature Backbone.** Given an input video-text/audio pair, we first sample  $T$  frames from the video clip, and utilize the visual backbone and a pre-trained text/audio backbone to extract the image and multi-modal features. Specifically, we utilize ResNet (He et al. 2016) or Swin Transformer (Liu et al. 2021) as the visual backbone, and obtain the multi-scale visual features of the  $2^{nd}$ ,  $3^{rd}$ ,  $4^{th}$  stages. Concurrently, for the text reference, we employ an off-the-shelf language model, RoBERTa (Liu et al. 2019), to encode the linguistic embedding tokens. For the audio reference, we first process it as a spectrogram transform via a short-time Fourier Transform and then feed it into a pre-trained VGGish (Hershey et al. 2017) model. After the text/audio encoding, a linear projection layer is adopted to align the multi-modal feature dimension with the visual features. Note that, following previous work (Wu et al. 2022), we adopt an early fusion module in the visual backbone to inject preliminary text/audio knowledge into visual features.

**MTA Module.** On top of feature extraction, we feed the visual and text/audio features into the multi-scale temporal

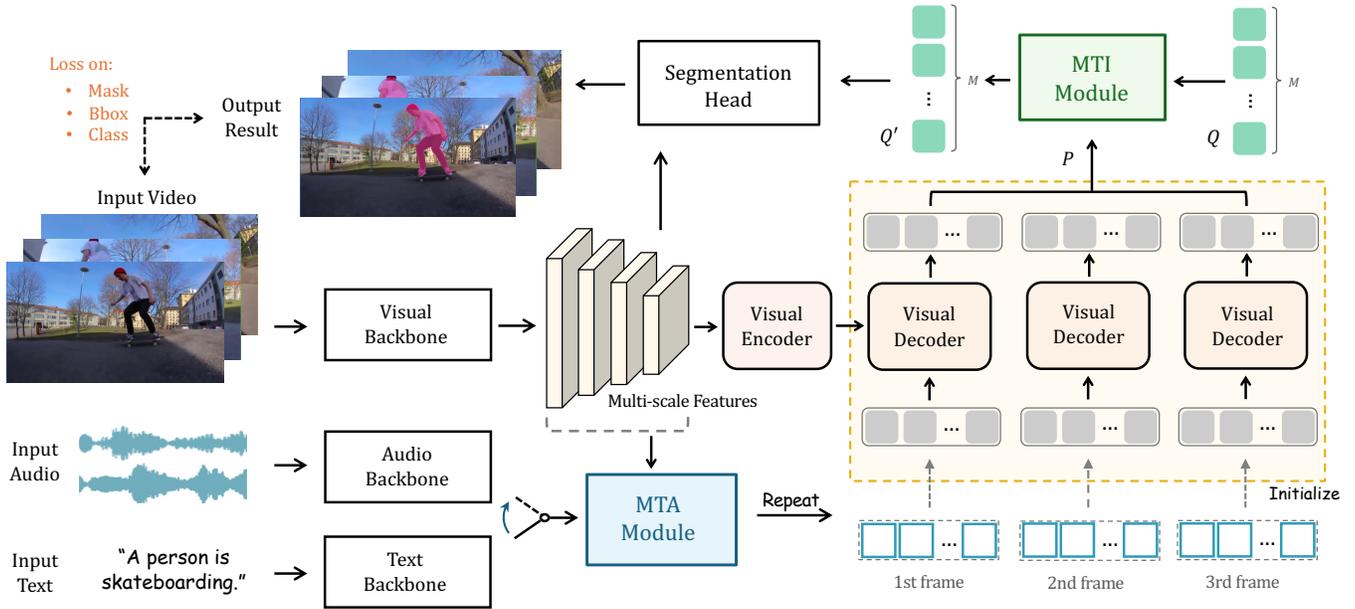


Figure 1: The Overall Pipeline of MUTR for referring video object segmentation. We present a unified transformer architecture to tackle video object segmentation referred by multi-modal inputs. We propose MTA module and MTI module for low-level multi-scale aggregation and high-level multi-object interaction, respectively.

aggregation module (MTA). We concatenate the visual features of adjacent frames, and adopt cascaded cross-attention blocks to enhance the multi-scale and multi-modal feature fusion, which is specifically described in Section .

**Visual Encoder-decoder Transformer.** The basic transformer consists of a visual encoder and a visual decoder, which processes the video in a frame-independent manner to focus on the feature fusion within a single frame. In detail, the visual encoder adopts vanilla self-attention blocks to encode the multi-scale visual features. The visual decoder regards the encoded visual features as the key and value, and the output references from the MTA module as learnable object queries for decoding. Unlike the randomly initialized queries in traditional DETR (Carion et al. 2020), ours are input-conditioned ones obtained via MTA module, which contains video-level multi-modal prior knowledge. With the visual decoder, the object queries gain rich instance information, which provides effective cues for the final segmentation process.

**MTI Module.** After the visual transformer, a multi-object temporal interaction (MTI) module is proposed for object-wise interaction, which is described in Section . In detail, we utilize an MTI encoder to communicate temporal features of the same object in different views. Then an MTI decoder is proposed to grasp information into a set of video-wise query representations for associating objects across frames, inspired by (Heo et al. 2022).

**Segmentation Head and Loss Function.** On top of the components introduced above, we obtain the final mask predictions from the extracted multi-modal features via a segmentation head. We follow previous works (Wu et al. 2022,

2021) to design the segmentation head that contains a bounding box head, a classification head, and a mask head. Then, we find the best assignment from the predictions of MUTR by using Hungarian Matching (Carion et al. 2020). During training, we calculate three losses in MUTR, which are focal loss (Lin et al. 2017)  $\mathcal{L}_{cls}$  on the predictions of referred object sequence,  $\mathcal{L}_{box}$  on the bounding box of predicted instance, and  $\mathcal{L}_{mask}$  on the predicted object masks. In detail,  $\mathcal{L}_{box}$  is the combination of  $L_1$  loss and GIoU loss (Rezatofighi et al. 2019), and  $\mathcal{L}_{mask}$  is the summation of the Dice (Milletari, Navab, and Ahmadi 2016) and binary focal loss. The whole loss function is formulated as

$$\mathcal{L} = \lambda_{cls} \mathcal{L}_{cls} + \lambda_{box} \mathcal{L}_{box} + \lambda_{mask} \mathcal{L}_{mask}, \quad (1)$$

where  $\lambda_{cls}$ ,  $\lambda_{box}$  and  $\lambda_{mask}$  denote the weights for  $\mathcal{L}_{cls}$ ,  $\mathcal{L}_{box}$  and  $\mathcal{L}_{mask}$ .

### Multi-scale Temporal Aggregation

To boost both the multi-modal and multi-frame feature fusion, we introduce **Multi-scale Temporal Aggregation** module for low-level temporal aggregation. The proposed MTA module generates a set of object queries that contain multi-modal knowledge for subsequent transformer decoding.

**Multi-scale Temporal Transform.** As shown in Figure ??, the MTA module take the text/audio features  $F_r$ , and multi-scale visual features as input, i.e., the extracted features of  $2^{nd}$ ,  $3^{rd}$ ,  $4^{th}$  stages from the visual backbone. We first utilize linear projection layers on the multi-scale features to transform them into the same dimension. Specifically, we separately utilize  $1 \times 1$  convolution layers on the  $2^{nd}$ ,  $3^{rd}$ ,  $4^{th}$

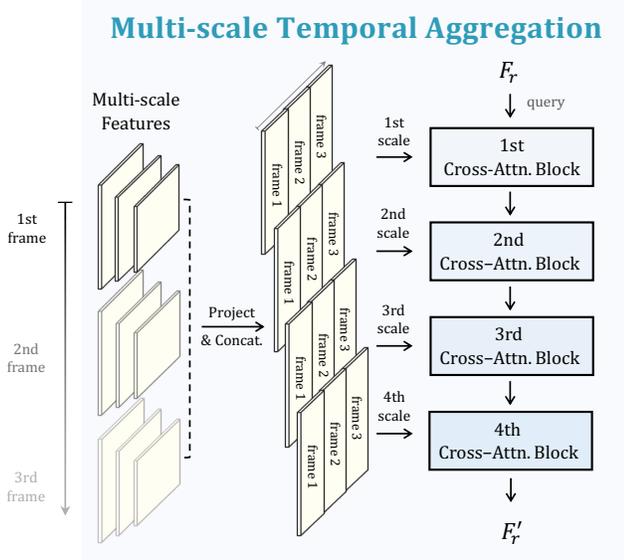


Figure 2: Multi-scale Temporal Aggregation. For low-level multi-modal temporal aggregation, we propose MTA module for inter-frame interaction, which generates tokens with multi-modal knowledge as the input queries for transformer decoding.

scale features, and an additional  $3 \times 3$  convolution layer on the  $4^{th}$  stage features to obtain the  $5^{th}$  scale features. We denote the projected features as  $\{F_{vj}^i\}$ , where  $2 \leq i \leq 5$ ,  $1 \leq j \leq T$  represent the stage number and frame number. After that, we concatenate the visual features of adjacent frames for each scale, formulated as

$$F_v^i = \text{Concat}(F_{v1}^i, F_{v2}^i, \dots, F_{vj}^i, \dots, F_{vT}^i), \quad (2)$$

where  $2 \leq i \leq 5$ ,  $1 \leq j \leq T$ ,  $F_{vj}^i$  represents the projected  $j^{th}$  frame features of  $i^{th}$  scale, and  $\{F_{vj}^i\}_{i=2}^5$  is the final transformed multi-scale visual feature. Then, the resulting multi-modal temporal features are regarded as the key and value in the following cross-attention blocks.

**Multi-modal Cross-attention.** On top of this, we adopt sequential cross-attention mechanisms for multi-modal tokens to progressively capture temporal visual cues of different image scales. We adopt four cross-attention blocks that are assigned to each scale respectively for multi-scale temporal feature extracting. In each attention block, the text/audio features serve as the query, while the multi-scale visual features serve as the key and value. We formulate it as

$$F_f = \text{Block}_{i-1}(F_r, F_v^i, F_v^i), \quad 2 \leq i \leq 5, \quad (3)$$

where  $\text{Block}$  represents the sequential cross-attention blocks in MTA module,  $F_f$  is the output multi-modal tokens that contain the multi-modal information.

After that, we simply repeat the class token of  $F_f$  for  $T \times N$  times, where  $T$  is the frame number and  $N$  is the query number. We adopt them as the initialized queries fed into the visual transformer for frame-wise decoding. With the

## Multi-object Temporal Interaction

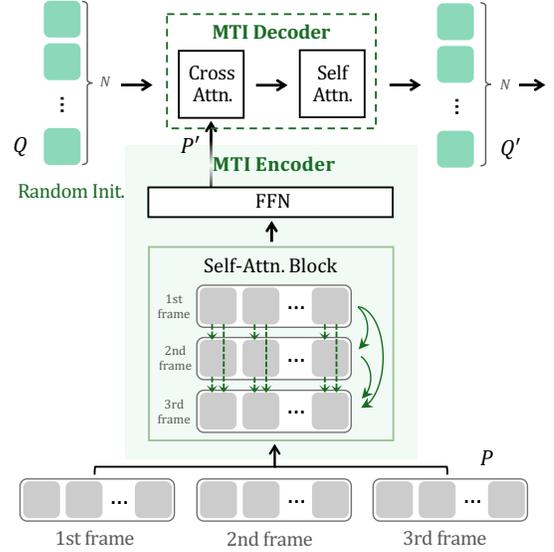


Figure 3: Multi-object Temporal Interaction. We introduce MTI module for inter-frame object-wise interaction, and maintain a set of video-wise query representations for associating objects across frames.

MTA module, the pre-initialized input queries obtain prior multi-scale knowledge and temporal information for better multi-modal alignment during subsequent decoding.

## Multi-object Temporal Interaction

As the visual transformer adopts a frame-independent manner and fails to interact information among multiple frames, we further introduce a **Multi-object Temporal Interaction** module to conduct inter-frame object-wise interaction. This module enhances the high-level temporal communication of objects, and benefits the visual correspondence for effective segmentation. The details of MTI are shown in Figure ??, which consists of an MTI encoder and an MTI decoder.

**MTI Encoder.** We obtain the object query outputs  $P$  of each frame from the transformer decoder, and feed them into the MTI encoder, which contains a self-attention layer to conduct object-wise interaction across multiple frames, and a feed-forward network layer for feature transformation. To achieve more efficient implementation, we adopt shifted window-attention (Liu et al. 2021) with linear computational complexity in the self-attention layer. The process of MTI encoder is formulated as

$$P' = \text{MTI\_Encoder}(P) \quad (4)$$

where  $\text{MTI\_Encoder}$  denotes the MTI encoder, and  $P'$  is the outputs of MTI encoder.

**MTI Decoder.** Based on the MTI encoder, we maintain a set of video-wise query  $Q$  for associating objects across frames, which are randomly initialized. We regard the outputs from MTI encoder as the key and value, and feed them and

Method	Backbone	Ref-YouTube-VOS			Ref-DAVIS 2017		
		$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
CMSA (Ye et al. 2019)	ResNet-50	34.9	33.3	36.5	34.7	32.2	37.2
URVOS (Seo, Lee, and Han 2020)		47.2	45.3	49.2	51.5	47.3	56.0
LBDT-4 (Ding et al. 2022b)		48.2	50.6	49.4	-	-	-
YOFO (Li et al. 2022)		48.6	47.5	49.7	53.3	48.8	57.9
ReferFormer (Wu et al. 2022)		58.7	57.4	60.1	61.1	58.0	64.1
<b>MUTR</b>		<b>61.9</b>	<b>60.4</b>	<b>63.4</b>	<b>65.3</b>	<b>62.4</b>	<b>68.2</b>
CITD (Liang et al. 2021)	ResNet-101	56.4	54.8	58.1	-	-	-
ReferFormer (Wu et al. 2022)		59.3	58.1	60.4	61.0	58.1	63.8
<b>MUTR</b>		<b>63.6</b>	<b>61.8</b>	<b>65.4</b>	<b>65.3</b>	<b>61.9</b>	<b>68.6</b>
ReferFormer (Wu et al. 2022)	Swin-L	64.2	62.3	66.2	63.9	60.8	67.0
<b>MUTR</b>		<b>68.4</b>	<b>66.4</b>	<b>70.4</b>	<b>68.0</b>	<b>64.8</b>	<b>71.3</b>
MTTR (Botach, Zheltonozhskii, and Baskin 2022)	Video-Swin-T	55.3	54.0	56.6	-	-	-
MANet (Chen et al. 2022)		55.6	54.8	56.5	-	-	-
ReferFormer (Wu et al. 2022)		62.6	59.9	63.3	62.8	60.8	67.0
<b>MUTR</b>		<b>64.0</b>	<b>62.2</b>	<b>65.8</b>	<b>66.5</b>	<b>63.0</b>	<b>70.0</b>
VLT (Ding et al. 2022a)	Video-Swin-B	63.8	61.9	65.6	61.6	58.9	64.3
ReferFormer (Wu et al. 2022)		64.9	62.8	67.0	64.3	60.7	68.0
<b>MUTR</b>		<b>67.5</b>	<b>65.4</b>	<b>69.6</b>	<b>66.4</b>	<b>62.8</b>	<b>70.0</b>

Table 1: Performance of MUTR on Ref-YouTube-VOS and Ref-DAVIS 2017 Datasets. We report the results of MUTR and prior works on multiple backbones, where our MUTR shows the *state-of-the-art* performance on all datasets.

video-wise queries  $Q$  into MTI decoder for video-wise decoding. The MTI decoder consists of a cross-attention layer, a self-attention layer, and a feed-forward network layer. We formulate them as

$$Q' = \text{MTI\_Decoder}(Q, P', P') \quad (5)$$

where  $\text{MTI\_Decoder}$  represents the MTI decoder,  $Q'$  is the outputs of MTI decoder. In this way, the proposed MTI module promotes high-level temporal fusion and enhances the connection and interaction of the same objects in different frames, which further contributes to effective segmentation.

### Joint Training for Multi-modality

As a unified VOS framework for multi-modality, MUTR has the potential to segment video objects referred by either text or audio reference. To achieve this, we conduct joint training by combining both text- and audio-referred datasets. Specifically, to balance the data amount of two modalities, the joint training data is composed of partial Ref-YouTube-VOS (Seo, Lee, and Han 2020) (text reference) and the entire AVSBench S4 (Zhou et al. 2022) (audio reference). We sample a subset of Ref-YouTube-VOS for training (10,093 clips (5 frames per clip) out of 72,920), for which we utilize only one description for videos with multiple text descriptions, and filter out half of the instances based on odd-index positions for training.

For text or audio reference, we accordingly switch to their respective encoders for feature encoding, i.e., RoBERTa for text and VGGish for audio. Then, they share the same subsequent network modules, i.e., MTA, visual encoder, visual decoder, MTI, and segments head. By our temporal and cross-modality interaction modules, the jointly trained MUTR can obtain superior performance on either of the two modalities.

## Experiments

### Quantitative Results

**Ref-YouTube-VOS.** As shown in Table 1, MUTR outperforms the previous state-of-the-art methods by a large margin under on all datasets. On Ref-YouTube-VOS, MUTR with a lightweight backbone ResNet-50 achieves the superior performance with overall  $\mathcal{J}\&\mathcal{F}$  of 61.9%, an improvement of +3.2% than the previous state-of-the-art method ReferFormer. By adopting a more powerful backbone Swin-Transformer (Liu et al. 2021), MUTR improves the performance to  $\mathcal{J}\&\mathcal{F}$  68.4%, which is +4.2% than the previous method ReferFormer (Wu et al. 2022). Using a more strong backbone, our method has a higher percentage of improvement, which better reflects the robustness of our method on the scaled-up model size. To reflect the powerful temporal modeling capability of MUTR, we therefore adopt the video Swin transformer (Liu et al. 2022) as the backbone, which is a spatial-temporal encoder that can effectively capture the spatial and temporal cues simultaneously, to compensate for the temporal limitations of the ReferFormer as discussed in (Hu et al. 2022). It can be observed that our method significantly outperforms the ReferFormer, which demonstrates the effectiveness of the temporal consistency in our model.

**Ref-DAVIS 2017.** On the Ref-DAVIS 2017, our method also achieves the best results under the same backbone setting. Since ReferFormer (Wu et al. 2022) does not include the resultson Ref-DAVIS 2017, we report its results using the official pre-trained models provided by ReferFormer.

**AV-VOS.** Table 2 shows the performance of our MUTR on the AVSBench dataset. MUTR significantly surpasses all the

Method	Backbone	AVSBench S4			AVSBench MS3		
		$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
LVS (Chen et al. 2021)	ResNet-18	44.5	37.9	51.0	31.3	29.5	33.0
SST (Duke et al. 2021)	ResNet-50	73.2	66.3	80.1	49.9	42.6	57.2
LGVT (Zhang et al. 2021)	Swin-B	81.1	74.9	87.3	50.0	40.7	59.3
Baseline (Zhou et al. 2022)	ResNet-50	78.8	72.8	84.8	52.9	47.9	57.8
Baseline (Zhou et al. 2022)	PvT-V2	83.3	78.7	87.9	59.3	54.0	64.5
MUTR	ResNet-50	83.0	78.6	87.3	61.6	57.0	66.1
	ResNet-101	83.1	78.5	87.6	63.7	59.0	68.3
	PvT-V2	85.1	80.7	89.5	67.9	63.7	72.0
	Swin-L	85.7	81.5	89.8	69.0	65.0	73.0
	Video-Swin-T	83.0	78.7	87.2	64.0	59.2	68.7
	Video-Swin-S	84.1	79.8	88.3	67.3	62.7	71.8
	Video-Swin-B	85.7	81.6	89.7	68.8	64.0	73.5

Table 2: Performance of MUTR on AVSBench Dataset. MUTR surpasses the *state-of-the-art* method.

Methods	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
ReferFormer (Wu et al. 2022)	32.5	32.6	32.4
<b>MUTR*</b>	39.9	39.4	40.5
<b>MUTR</b>	41.3	40.6	42.0

Table 3: Performance of MUTR on Ref-YouTube-VOS by Multi-modality Joint Training.

Methods	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
Baseline (Zhou et al. 2022)	78.8	72.8	84.8
<b>MUTR*</b>	79.7	74.5	84.9
<b>MUTR</b>	81.4	76.8	85.9

Table 4: Performance of MUTR on AVSBench S4 by Multi-modality Joint Training.

Components		Block Num.	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
Multi-scale	Temporal				
✓	-	1	61.3	59.7	62.7
-	✓	1	60.4	58.9	61.9
✓	✓	1	<b>61.9</b>	<b>60.4</b>	<b>63.4</b>
✓	✓	2	60.7	59.3	62.2
✓	✓	3	60.4	59.1	61.7

Table 5: Ablation Study of MTA Module.

Components		Block Num.	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$
Encoder	Decoder				
✓	-	3	60.3	58.8	61.9
-	✓	3	61.2	60.0	62.6
✓	✓	3	<b>61.9</b>	<b>60.4</b>	<b>63.4</b>
✓	✓	2	61.1	59.5	62.6
✓	✓	1	60.8	59.3	62.3

Table 6: Ablation Study of MTI Module.

previous best competitors ( $\mathcal{J}\&\mathcal{F}$  **83.0% VS 78.8%**; **61.6% VS 52.9%**) with the same ResNet-50 backbone. We also achieve a new state-of-the-art performance with Swin-L (Liu et al. 2021) backbone. By employing a stronger backbone, we observe consistent performance improvement of MUTR, indicating the strong generalization of our approach.

**Joint Training Datasets.** We keep most training hyperparameters consistent with our previous text-referred video object segmentation experiments, and adopt ResNet-50 as the visual backbone. Table 3 and 4 present the performance of MUTR by joint training on Ref-YouTube-VOS and AVSBench S4, respectively. Therein, ReferFormer, the ‘Baseline’, and MUTR\* are all trained exclusively on text- or audio-referred dataset, while MUTR is trained on the multi-modality joint dataset. As shown, the single unified MUTR by joint training can achieve even better performance than their separate training. This indicates the effectiveness of our proposed architecture to serve as a unified framework

simultaneously for text and audio input.

## Qualitative Results

The first two columns of Figure 4 visualize some qualitative results in comparison with ReferFormer (Wu et al. 2022), which lacks inter-frame interaction in terms of temporal dimension. As demonstrated, along with multiple highly similar objects in the video, ReferFormer (Wu et al. 2022) is easier to misidentifies them. In contrast, MUTR can associate all the objects in temporal, which can better track and segment all targets accurately.

The last column of Figure 4 visualizes the audio-visual result compared with Baseline (Zhou et al. 2022) on AVSBench S4. With temporal consistency, MUTR successfully tracks and segments challenging situations that are surrounded or occluded by similar instances.

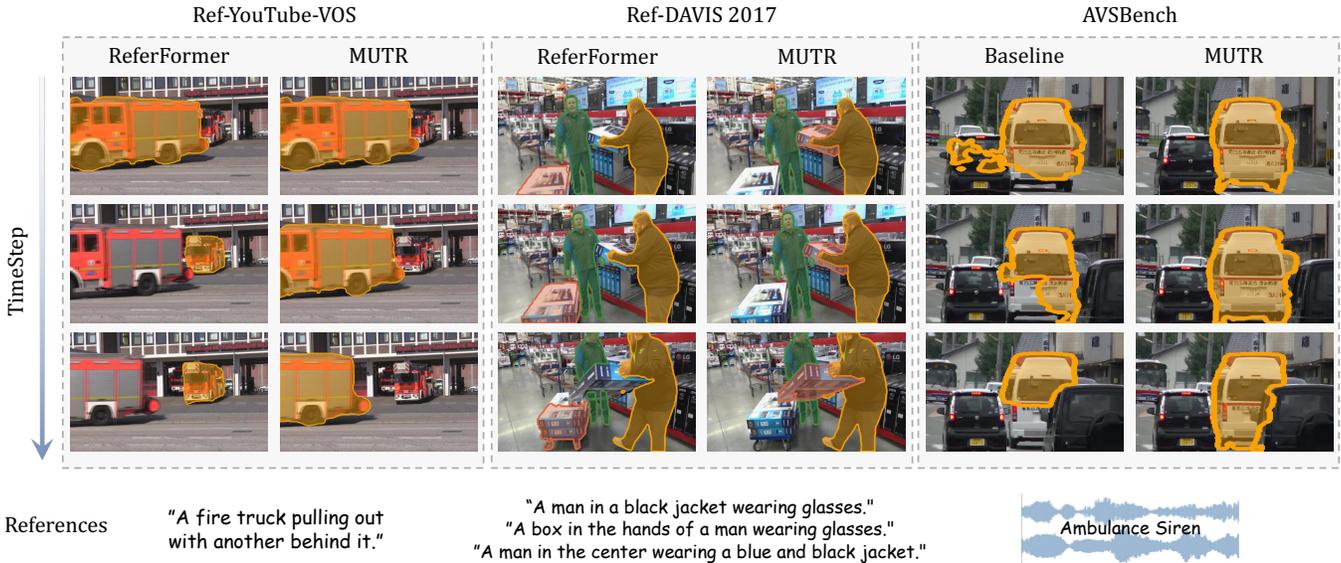


Figure 4: Qualitative Results of MUTR. We visualize the results between ReferFormer (Wu et al. 2022) and MUTR on R-VOS benchmarks and between Baseline (Zhou et al. 2022) and MUTR on AV-VOS benchmark. Compared with ReferFormer, MUTR performs better on temporal consistency when segmenting multiple similar objects, i.e., fire truck in Ref-YouTube-VOS and box in Ref-DAVIS 2017. Also, compared with the baseline of AV-VOS (Zhou et al. 2022) that denoted as ‘Baseline’ in this figure, MUTR can handle serve occlusion.

MTA	MTI	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}$	$\mathcal{F}$	FPS	Parameters
-	-	60.2	58.7	61.7	19.64	168.1M
-	✓	60.8	59.3	62.2	19.53	176.4M
✓	-	61.5	60.1	63.0	19.44	169.3M
✓	✓	<b>61.9</b>	<b>60.4</b>	<b>63.4</b>	19.37	177.6M

Table 7: Ablation Study of the MTA and MTI Modules.

### Ablation Studies

In this section, we perform experiments to analyze the main components and hyper-parameters of MUTR. All the experiments are conducted with the ResNet-50 backbone and evaluate their impact by the Ref-YouTube-VOS performance.

**Effectiveness of Main Componentenets.** Table 7 demonstrates the effectiveness of MTA and MTI proposed in our framework. The performance will be seriously degraded from 61.9% to 60.2% by removing MTA and MTI modules. Besides, our MTA and MTI modules introduce a marginal increase in inference latency, demonstrating favorable implementation and parameter efficiency.

**Ablation Study on MTA.** In Table 5, if either the single-scale temporal aggregation or multi-scale aggregation at the image level are adopted, the performance of MUTR would significantly drop to 60.4% and 61.3%, respectively, which demonstrates the necessity of the MTA module. We also ablate the number of MTA blocks. As seen in Table 5, more MTA blocks cannot bring further performance improvement, since (1) not enough videos for training; (2) the embedding space of visual and reference is only 256-dimensional, which is difficult to optimize so many parameters.

**Ablation Study on MTI.** As shown in Table 6, the performance of MUTR is improved by using more MTI blocks. A possible reason is that the larger the MTI blocks, the more sufficient temporal communication between instance-level can be performed. Moreover, using only the encoder or decoder, the performance of MUTR would both decline.

### Conclusion

This paper proposes a MUTR, a Multi-modal Unified Temporal transformer for Referring video object segmentation. A simple yet and effective Multi-scale Temporal Aggregation (MTA) is introduced for multi-modal references to explore low-level multi-scale visual information in video-level. Besides, the high-level Multi-object Temporal Interaction (MTI) is designed for inter-frame feature communication to achieve temporal correspondence between the instance-level across the entire video. Aided by the MTA and MTI, our MUTR achieves new state-of-the-art performance on three R-VOS/AV-VOS benchmarks compared to previous solutions. We hope the MTA and MTI will help ease the future study of multi-modal VOS and related tasks (e.g., referring video object tracking and video instance segmentation). We do not foresee negative social impact from the proposed work.

### Acknowledgements

This work was supported by National Natural Science Foundation of China (Grant No. 62206272). This work was also supported in part by Scientific and Technological Innovation Action Plan of Shanghai Science and Technology Committee (No.22511101502 and No.21DZ2203300).

## References

- Bellver, M.; Ventura, C.; Silberer, C.; Kazakos, I.; Torres, J.; and Giro-i Nieto, X. 2020. Refvos: A closer look at referring expressions for video object segmentation. *arXiv preprint arXiv:2010.00263*.
- Botach, A.; Zheltonozhskii, E.; and Baskin, C. 2022. End-to-end referring video object segmentation with multimodal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4985–4995.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, 213–229. Springer.
- Chen, D.-J.; Jia, S.; Lo, Y.-C.; Chen, H.-T.; and Liu, T.-L. 2019. See-through-text grouping for referring image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7454–7463.
- Chen, H.; Xie, W.; Afouras, T.; Nagrani, A.; Vedaldi, A.; and Zisserman, A. 2021. Localizing visual sounds the hard way. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16867–16876.
- Chen, W.; Hong, D.; Qi, Y.; Han, Z.; Wang, S.; Qing, L.; Huang, Q.; and Li, G. 2022. Multi-Attention Network for Compressed Video Referring Object Segmentation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4416–4425.
- Ding, H.; Liu, C.; Wang, S.; and Jiang, X. 2021. Vision-language transformer and query generation for referring segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16321–16330.
- Ding, H.; Liu, C.; Wang, S.; and Jiang, X. 2022a. VLT: Vision-Language Transformer and Query Generation for Referring Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ding, Z.; Hui, T.; Huang, J.; Wei, X.; Han, J.; and Liu, S. 2022b. Language-bridged spatial-temporal interaction for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4964–4973.
- Duke, B.; Ahmed, A.; Wolf, C.; Aarabi, P.; and Taylor, G. W. 2021. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5912–5921.
- Fang, R.; Yan, S.; Huang, Z.; Zhou, J.; Tian, H.; Dai, J.; and Li, H. 2023. InstructSeq: Unifying Vision Tasks with Instruction-conditioned Multi-modal Sequence Generation. *arXiv preprint arXiv:2311.18835*.
- Gao, P.; Han, J.; Zhang, R.; Lin, Z.; Geng, S.; Zhou, A.; Zhang, W.; Lu, P.; He, C.; Yue, X.; et al. 2023. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*.
- Guo, Z.; Zhang, R.; Zhu, X.; Tang, Y.; Ma, X.; Han, J.; Chen, K.; Gao, P.; Li, X.; Li, H.; et al. 2023. Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint arXiv:2309.00615*.
- Han, J.; Gong, K.; Zhang, Y.; Wang, J.; Zhang, K.; Lin, D.; Qiao, Y.; Gao, P.; and Yue, X. 2023a. OneLLM: One Framework to Align All Modalities with Language. *arXiv preprint arXiv:2312.03700*.
- Han, J.; Zhang, R.; Shao, W.; Gao, P.; Xu, P.; Xiao, H.; Zhang, K.; Liu, C.; Wen, S.; Guo, Z.; et al. 2023b. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint arXiv:2309.03905*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heo, M.; Hwang, S.; Oh, S. W.; Lee, J.-Y.; and Kim, S. J. 2022. Vita: Video instance segmentation via object token association. *arXiv preprint arXiv:2206.04403*.
- Hershey, S.; Chaudhuri, S.; Ellis, D. P.; Gemmeke, J. F.; Jansen, A.; Moore, R. C.; Plakal, M.; Platt, D.; Saurous, R. A.; Seybold, B.; et al. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*, 131–135. IEEE.
- Hu, Z.; Chen, B.; Gao, Y.; Ji, Z.; and Bai, J. 2022. 1st Place Solution for YouTubeVOS Challenge 2022: Referring Video Object Segmentation. *arXiv preprint arXiv:2212.14679*.
- Hu, Z.; Feng, G.; Sun, J.; Zhang, L.; and Lu, H. 2020. Bi-directional relationship inferring network for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4424–4433.
- Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1780–1790.
- Khoreva, A.; Rohrbach, A.; and Schiele, B. 2019. Video object segmentation with language referring expressions. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part IV 14*, 123–141. Springer.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.
- Li, D.; Li, R.; Wang, L.; Wang, Y.; Qi, J.; Zhang, L.; Liu, T.; Xu, Q.; and Lu, H. 2022. You only infer once: Cross-modal meta-transfer for referring video object segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1297–1305.
- Liang, C.; Wu, Y.; Zhou, T.; Wang, W.; Yang, Z.; Wei, Y.; and Yang, Y. 2021. Rethinking cross-modal interaction from a top-down perspective for referring video object segmentation. *arXiv preprint arXiv:2106.01061*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.

- Lin, Z.; Liu, C.; Zhang, R.; Gao, P.; Qiu, L.; Xiao, H.; Qiu, H.; Lin, C.; Shao, W.; Chen, K.; et al. 2023. SPHINX: The Joint Mixing of Weights, Tasks, and Visual Embeddings for Multi-modal Large Language Models. *arXiv preprint arXiv:2311.07575*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu, Z.; Ning, J.; Cao, Y.; Wei, Y.; Zhang, Z.; Lin, S.; and Hu, H. 2022. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3202–3211.
- Luo, G.; Zhou, Y.; Sun, X.; Cao, L.; Wu, C.; Deng, C.; and Ji, R. 2020. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 10034–10043.
- Margffoy-Tuay, E.; Pérez, J. C.; Botero, E.; and Arbeláez, P. 2018. Dynamic multimodal instance segmentation guided by natural language queries. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 630–645.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, 565–571. Ieee.
- Mo, S.; and Tian, Y. 2023. AV-SAM: Segment Anything Model Meets Audio-Visual Localization and Segmentation. *arXiv preprint arXiv:2305.01836*.
- Rezatofghi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 658–666.
- Seo, S.; Lee, J.-Y.; and Han, B. 2020. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, 208–223. Springer.
- Shi, H.; Li, H.; Meng, F.; and Wu, Q. 2018. Key-word-aware network for referring expression image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 38–54.
- Wang, K.; Ren, H.; Zhou, A.; Lu, Z.; Luo, S.; Shi, W.; Zhang, R.; Song, L.; Zhan, M.; and Li, H. 2023. MathCoder: Seamless Code Integration in LLMs for Enhanced Mathematical Reasoning. *arXiv preprint arXiv:2310.03731*.
- Wang, Z.; Lu, Y.; Li, Q.; Tao, X.; Guo, Y.; Gong, M.; and Liu, T. 2022. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11686–11695.
- Wu, J.; Jiang, Y.; Bai, S.; Zhang, W.; and Bai, X. 2021. SeqFormer: Sequential Transformer for Video Instance Segmentation. *arXiv preprint arXiv:2112.08275*.
- Wu, J.; Jiang, Y.; Sun, P.; Yuan, Z.; and Luo, P. 2022. Language as queries for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4974–4984.
- Xu, N.; Yang, L.; Fan, Y.; Yue, D.; Liang, Y.; Yang, J.; and Huang, T. 2018. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*.
- Yan, S.; Xu, X.; Zhang, R.; Hong, L.; Chen, W.; Zhang, W.; and Zhang, W. 2023. PanoVOS: Bridging Non-panoramic and Panoramic Views with Transformer for Video Segmentation. *arXiv e-prints*, arXiv–2309.
- Yang, Z.; Wang, J.; Tang, Y.; Chen, K.; Zhao, H.; and Torr, P. H. 2022. Lavt: Language-aware vision transformer for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18155–18165.
- Ye, L.; Rochan, M.; Liu, Z.; and Wang, Y. 2019. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10502–10511.
- Zhang, J.; Xie, J.; Barnes, N.; and Li, P. 2021. Learning generative vision transformer with energy-based latent space for saliency prediction. *Advances in Neural Information Processing Systems*, 34: 15448–15463.
- Zhang, R.; Han, J.; Zhou, A.; Hu, X.; Yan, S.; Lu, P.; Li, H.; Gao, P.; and Qiao, Y. 2023a. LLaMA-Adapter: Efficient Fine-tuning of Language Models with Zero-init Attention. *arXiv preprint arXiv:2303.16199*.
- Zhang, R.; Hu, X.; Li, B.; Huang, S.; Deng, H.; Li, H.; Qiao, Y.; and Gao, P. 2023b. Prompt, Generate, then Cache: Cascade of Foundation Models makes Strong Few-shot Learners. *CVPR 2023*.
- Zhang, R.; Jiang, Z.; Guo, Z.; Yan, S.; Pan, J.; Dong, H.; Gao, P.; and Li, H. 2023c. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*.
- Zhang, R.; Zhang, W.; Fang, R.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2022. Tip-Adapter: Training-free Adaption of CLIP for Few-shot Classification. In *ECCV 2022*. Springer Nature Switzerland.
- Zhou, J.; Wang, J.; Zhang, J.; Sun, W.; Zhang, J.; Birchfield, S.; Guo, D.; Kong, L.; Wang, M.; and Zhong, Y. 2022. Audio-Visual Segmentation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII*, 386–403. Springer.
- Zhu, X.; Zhang, R.; He, B.; Zhou, A.; Wang, D.; Zhao, B.; and Gao, P. 2023. Not all features matter: Enhancing few-shot clip with adaptive prior refinement. *ICCV 2023*.