

# Cooperation Does Matter: Exploring Multi-Order Bilateral Relations for Audio-Visual Segmentation

Qi Yang<sup>1,2</sup> Xing Nie<sup>1,2</sup> Tong Li<sup>3</sup> Pengfei Gao<sup>3</sup> Ying Guo<sup>3</sup>  
 Cheng Zhen<sup>3</sup> Pengfei Yan<sup>3</sup> Shiming Xiang<sup>1,2</sup>

<sup>1</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences (UCAS)

<sup>2</sup> Institute of Automation, Chinese Academy of Sciences (CASIA) <sup>3</sup> Meituan

## Abstract

Recently, an audio-visual segmentation (AVS) task has been introduced, aiming to group pixels with sounding objects within a given video. This task necessitates a first-ever audio-driven pixel-level understanding of the scene, posing significant challenges. In this paper, we propose an innovative audio-visual transformer framework, termed COMBO, an acronym for COoperation of Multi-order Bilateral relations. For the first time, our framework explores three types of bilateral entanglements within AVS: pixel entanglement, modality entanglement, and temporal entanglement. Regarding pixel entanglement, we employ a Siam-Encoder Module (SEM) that leverages prior knowledge to generate more precise visual features from the foundational model. For modality entanglement, we design a Bilateral-Fusion Module (BFM), enabling COMBO to align corresponding visual and auditory signals bi-directionally. As for temporal entanglement, we introduce an innovative adaptive inter-frame consistency loss according to the inherent rules of temporal. Comprehensive experiments and ablation studies on AVSBench-object (84.7 mIoU on S4, 59.2 mIoU on MS3) and AVSBench-semantic (42.1 mIoU on AVSS) datasets demonstrate that COMBO surpasses previous state-of-the-art methods. Project page is available at <https://yannqi.github.io/AVS-COMBO>.

## 1. Introduction

Human visual attention is often *hear-guided*, i.e., we tend to focus on the object with sounds [4]. For example, when we hear a cat meow, we pay more attention to the cat than other objects due to the strong association between the meow and the cat. Inspired by this potential interaction of auditory and visual signals, the cross-modal studies of vision and hearing have attracted numerous researchers, such as the audio-visual correspondence [1, 2], which only aims to match visual images and audio signals to the same

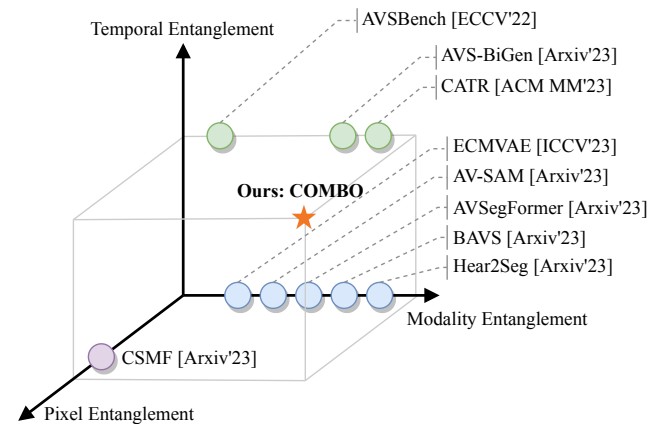


Figure 1. Comparison between the proposed COMBO and existing state-of-the-art methods. Our COMBO is the first work to simultaneously explore multi-order bilateral relations in modality, temporal and pixel levels.

scene, and sound source localization [5, 18, 34] which further seeks to locate the vocal visible regions. However, they have only focused on audio-visual tasks at the image or region levels, lacking pixel-level annotations. Recently, AVSBench [44, 45] integrates audio signals into video segmentation, called Audio-Visual Segmentation (AVS), which comprises two benchmarks: 1) AVSBench-object, which includes single source sound segmentation (S4) and multiple sound sources segmentation (MS3); 2) AVSBench-semantic, which further extends audio-visual semantic segmentation (AVSS) based on AVSBench-object.

Given that audio-visual segmentation is a burgeoning field that spans both audio and visual modalities, it presents a non-trivial task. Generally, when performing AVS, a cross-modal segmentation task involving video, there are mainly three challenges: (1) AVS contains audio and visual modalities, thus demanding explicit alignment of sequential audio features to spatial pixel-level activations; (2) AVS involves temporal information where the state of the current frame is dynamically affected by historical frames, therefore, exploring the correlation between adjacent frames is

essential; (3) AVS includes an image segmentation task; compared to 1D audio signals, 2D image signals have more redundant information, which is prone to be affected by background noise, thus requiring precise extraction of features from the image. To resolve the issue (1), the prevailing methods [13, 25, 26, 31, 33] employ matrix multiplication and modified cross-attention module to encode pixel-wise audio-visual interaction. Though impressive, these designs ignore the temporal dependence of adjacent frames that have been proven to be important for AVS. To solve this problem, some methods [15, 23, 44] partition temporal relations into consideration to explore both issues (1) and (2) simultaneously since AVS is a cross-modal video task. Nevertheless, their approaches rely too much on implicit inter-frame relations, leading to inaccurate associations. Regarding issue (3), CSMF [3] leverages frozen large foundation models to extract pure visual features for AVS. However, it independently tackles the audio and visual signals by naively combining several existing foundation models, resulting in sub-optimal performance.

To this end, we present **COMBO**, a novel audio-visual transformer framework for AVS. According to the three issues mentioned above, COMBO simultaneously considers modality, temporal, and pixel levels by introducing their bilateral entanglements, as shown in Fig. 1. Specifically, nature itself has many bilateral relations. For example, in electricity and magnetism, due to their intrinsic correlation, the change of current causes the change of magnetic field, and vice versa. Motivated by this, we refer to this bilateral relationship of mutual influence as entanglement.

In this work, we explore three potential bilateral entanglements: pixel entanglement, modality entanglement, and temporal entanglement. Pixel entanglement refers to the interdependent relationship between an image and its corresponding mask. Since background noise in the image leads to inaccuracies in the image-to-mask prediction process, it is essential to utilize external masks from the foundation models to entangle the input image to assist the model. Therefore, we construct a Siam-Encoder Module (SEM) as a visual feature extractor to facilitate more precise visual features, which can liberate from the constraints of foundation models than other methods [26, 33, 37]. Besides, as for the alignments of the audio and visual signals, we explore the modality entanglement between audio and visual components to amplify the efficiency of cross-modal matching. Contrary to existing single-fusion methods [13, 44], we believe that the cooperation between the two modalities can produce a positive effect. Inspired by [24], we initially propose a potent and memory-efficient bidirectional audio-visual fusion module called Bilateral-Fusion Module (BFM). Our BFM amplifies the spatial awareness of visual features relevant to sounding objects and strengthens the attention of audio signals embodying visual targets.

Moreover, the audio-visual tasks contain a solid temporal entanglement. Thus, we design an adaptive inter-frame consistency loss to better harness this inherent characteristic.

Our main contributions can be summarized as follows:

- We propose a Siam-Encoder Module (SEM) that transfers the knowledge of the foundation model for mining the potential pixel entanglement.
- We propose a Bilateral-Fusion Module (BFM) to take full advantage of the potential of both audio and visual modalities by exploring the modality entanglement.
- We propose an adaptive inter-frame consistency loss based on the inherent coherence of audio-visual tasks for enhanced temporal entanglement.
- We show that COMBO significantly outperforms existing state-of-the-art approaches in the challenging AVSBench-object and AVSBench-semantic datasets.

## 2. Related Work

### 2.1. Sound Source Localization

Sound source localization aims to estimate the position of a sound source in a video sequence, which is the most related task to the audio-visual segmentation task. LVS [5] utilizes a hard-mining strategy and a contrastive learning mechanism to discriminate challenging image fragments. DSOL [18] executes class-aware sounding object localization from mixed sound, which initially focuses on learning robust object representations from single-source localization. MSSL [34] localizes multiple sound sources in unconstrained videos without pairwise sound-object annotations. This approach involves the development of a two-stage learning framework, followed by the execution of cross-modal feature alignment. The pioneering methods in these areas have significantly inspired our research on AVS.

### 2.2. Semantic Segmentation

Semantic segmentation is a fundamental task that requires pixel-level classification. Early researchers take the Fully Convolutional Networks (FCN) [27] as the dominant approach and focus on aggregating long-range context in the feature map. PSPNet [42] performs spatial pyramid pooling at several grid scales. DeepLab [7, 8] utilizes atrous convolutions with different atrous rates. Furthermore, some methods [35, 38, 43] replace traditional convolutional backbone with transformer-based architectures. MaskFormer [9] and Mask2Former [11] propose a mask classifier with learnable queries and specialized designs for mask prediction. OneFormer [19] presents a universal image segmentation framework that unifies segmentation with a multi-task train-once design. Recently, a series of SAM models [20–22, 40] propose to build a foundation model for promptable segmentation with strong generalization. Given that the AVS task entails segmentation, these studies have significantly

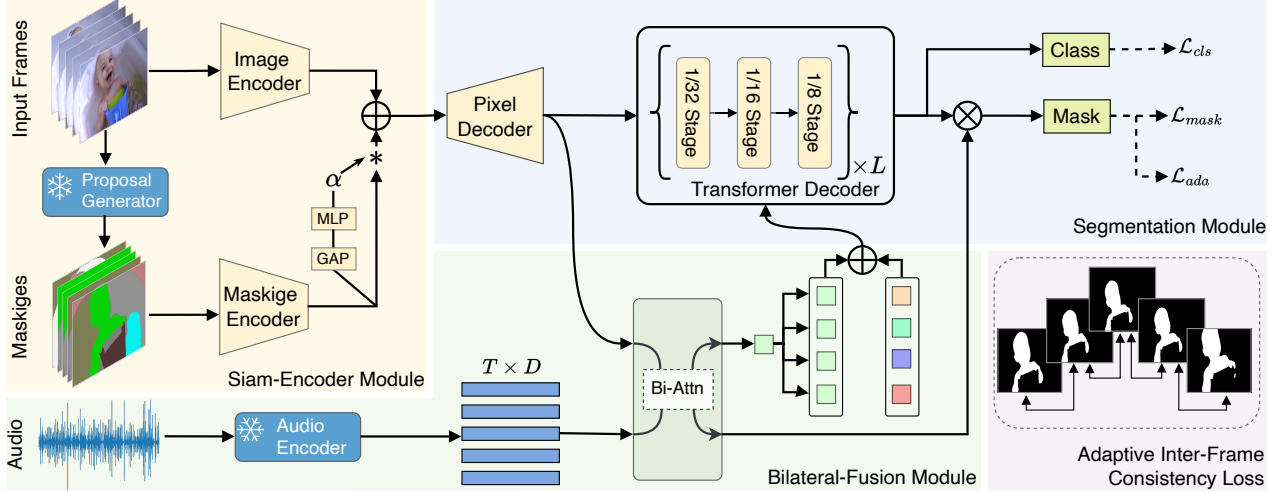


Figure 2. Overview of the proposed COMBO. COMBO adopts a novel audio-visual transformer framework specifically for audio-visual segmentation. Aiming at multi-order bilateral entanglement, our method is composed of three independent modules. (1) We introduce the Siam-Encoder Module, which is designed for the exploration of pixel entanglement. (2) To integrate the entanglement of audio and visual signals, we propose a Bilateral-Fusion Module. (3) Given the inherent characteristics of temporal entanglement, we construct an adaptive inter-frame consistency loss in the segmentation module to enhance the consistency of the output.

contributed to our work.

### 2.3. Audio-Visual Segmentation

AVS is an emerging task that aims to locate sounding sources by predicting pixel-wise maps and attracts many researchers [13, 23, 26, 30, 31, 33, 37, 45]. AVSBench [45] first constructs the audio-visual segmentation benchmark and proposes a temporal pixel-wise audio-visual interaction module (TPAVI) to inject audio semantics as guidance for the visual segmentation process. AVSegformer [13] proposes a transformer architecture that introduces audio features into the transformer decoder, enabling the network to attend to interested visual features selectively. CATR [23] proposes a combinatorial dependence fusion approach that comprehensively accounts for the spatial-temporal dependencies of audio-visual combination. Some methods [15, 30, 31] take advantage of the generative manners with latent diffusion model or variational auto-encoder to address AVS task. In addition, AV-SAM [33], GAVS [37], and BAVS [26] utilize the large foundation model to bootstrap audio-visual segmentation. Different from the above methods [13, 26, 37, 44], our proposed COMBO rethinks AVS from bilateral relations of three entanglements, which enhances the model’s representation ability by exploring the pixel, modality and temporal inherent relationships.

## 3. Method

### 3.1. Bilateral Visual Features Extraction

As illustrated in Fig. 2, our method initiates with the extraction of visual features primarily because the audio-visual segmentation (AVS), as a dense prediction task, ex-

hibits extensive pixel entanglement in visual perception. Recent studies [21, 22, 40] have demonstrated that the Segment Anything Model [20] exhibits robust generalization performance in segmentation tasks. Consequently, transferring the impressive capabilities of the foundation model to more complex visual tasks, such as AVS, presents an intriguing and valuable research question. The extension, however, is not straightforward. Although some methods [3, 33] attempt to fine-tune or concatenate the pre-trained SAM model for AVS, the limited capacity of the frozen foundation model restricts its performance to address the AVS task. Additionally, the AVS task aims to predict all sound targets per pixel, whereas the SAM model is only capable of generating class-agnostic masks without any audio guidance, thus demonstrating a significant disparity. Therefore, transferring the foundation model’s knowledge to the AVS task presents a tough challenge.

**Maskige as Prior Knowledge.** To solve the aforementioned issues, we believe that a feasible strategy is to incorporate the knowledge of the foundation model into visual features as a pixel entanglement, which is memory-efficient. Specifically, as depicted in the left area of Fig. 2, we introduce a Maskige generator. Given the input image  $x \in \mathbb{R}^{3 \times H \times W}$  and the frozen foundation model, one can derive the class-agnostic masks  $c \in \mathbb{R}^{K \times H \times W}$ , where  $K$  represents the number of potential targets. It is essential to highlight that  $K$  is dynamic and varies with input images. Thus, we first amplify the quantity of class-agnostic masks from  $K$  to  $N$  with zero masks and obtain a series of binary masks, where  $N$  is predetermined. However, given that the output of SAM is a series of binary masks, it is difficult to integrate them into visual features. Consequently,

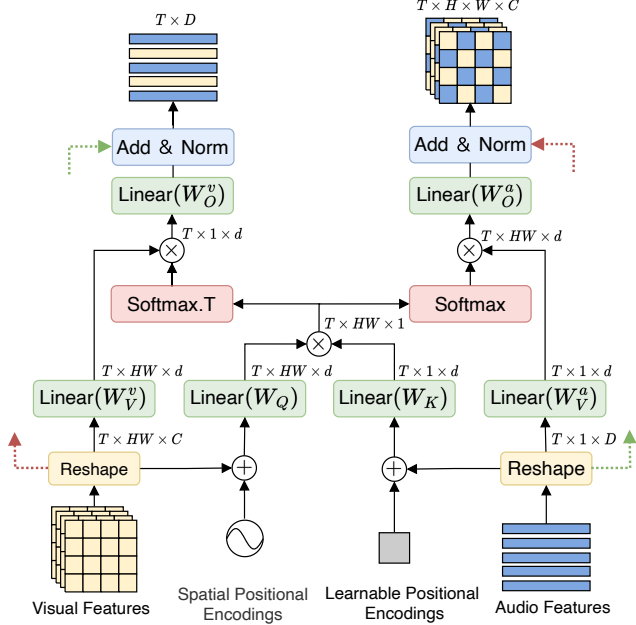


Figure 3. Illustration of Bilateral-Fusion Module (BFM). We input both visual and image signals, which are subsequently processed through bilateral attention to yield the fused visual and image features respectively. We omit the subscripts of  $H$  and  $W$  for better understanding. For enhanced visibility, the dashed line indicates a skip connection. Best viewed in color.

inspired by [6], we introduce Maskige  $m \in \mathbb{R}^{3 \times H \times W}$ , which shares the exact dimensions as the input image, to integrate prior knowledge better. To efficiently generate Maskige  $m$ , we employ a random color encoding function  $\mathcal{X}(\cdot) : \mathbb{R}^N \rightarrow \mathbb{R}^3$  that is capable of transforming the binary masks  $c \in \mathbb{R}^{N \times H \times W}$  into Maskige  $m \in \mathbb{R}^{3 \times H \times W}$  without extra training. Specifically,  $\mathcal{X}$  is designed to enhance the distinguishability of the Maskige and can be regarded as a linear layer  $\mathcal{X}(c) = cA$ , where  $A \in \mathbb{R}^{N \times 3}$ . To facilitate offline inference using the Maskige generator, the value of  $A$  is manually set appropriately without additional training. More details are in Appendix.

**Siam-Encoder Module (SEM).** To incorporate the image-like Maskiges as prior knowledge into input frames, we propose intertwining the features of Maskige and visual elements during the feature extraction stage. Accordingly, we design a Siam-Encoder Module, as depicted in Fig. 2. This module encompasses an Image encoder  $E_v$  and a Maskige encoder  $E_m$ , sharing a common framework. More precisely, for a short video clip with  $T$  frames  $I \in \mathbb{R}^{T \times 3 \times H \times W}$ , the Maskiges can be generated using the Maskige generator, resulting in  $M \in \mathbb{R}^{T \times 3 \times H \times W}$ . Subsequently, we extract multiple output features from both the image encoder and Maskige encoder, respectively. This process can be defined as follows:

$$F_{visual} = E_v(I; \theta_v), F_{visual} \in \{F_{v_i}\}_{i=1}^4, \quad (1)$$

$$F_{maskige} = E_m(M; \theta_m), F_{maskige} \in \{F_{m_i}\}_{i=1}^4, \quad (2)$$

in which  $F_{v_i}$  and  $F_{m_i} \in \mathbb{R}^{T \times \frac{H}{2^{i+1}} \times \frac{W}{2^{i+1}} \times C_i}$ .  $C_i$  represents the dimension of the  $i$ -th stage output features. To integrate the Maskige features into COMBO, we introduce channel-weighted blocks that augment the original visual features. The formula can be written as follows:

$$F_{v_i} = F_{m_i}(\text{GAP}(F_{m_i})W) + F_{v_i}, i = \{1, 2, 3, 4\}, \quad (3)$$

where GAP stands for global average pooling, and  $W \in \mathbb{R}^{C_i \times C_i}$  represents the linear weight. For simplicity, the bias is omitted in this context. After obtaining the Maskiges as prior information to boost pixel-level entanglement with visual features, the next critical aspect is exploring the modality entanglement between audio and visual signals.

### 3.2. Audio-Visual Bilateral Fusion

The relationship between any two modalities can be characterized as bilateral entanglement. For instance, an image can be described in text, and sound is inextricably linked with its visual counterpart. This entanglement among these distinct modalities provides an invaluable resource for researchers tackling multimodal tasks. Prior studies [38, 44] have overemphasized the influence of audio on visual features, thereby underestimating the significance of visual information to audio features. To address this imbalance, we propose a Bilateral-Fusion Module (BFM) in COMBO that surpasses a mere single fusion effect.

**Audio Feature Extraction.** For an audio clip corresponding to the input frames, we adopt VGGish [17] to extract audio features following [44]. Firstly, the audio clip is resampled to yield a 16kHz mono output  $A_{mono} \in \mathbb{R}^{N_{samples} \times 96 \times 64}$ , where  $N_{samples}$  is related to the duration of the audio. Then, a short-time Fourier transform is performed to yield a mel spectrum, denoted as  $A_{mel} \in \mathbb{R}^{T \times 96 \times 64}$ . Finally, the mel spectrum is subsequently fed into the VGGish model, resulting in the extraction of audio features  $F_a \in \mathbb{R}^{T \times D}$ , where  $T$  denotes the number of frames and  $D$  represents dimension of the audio.

**Bilateral-Fusion Module (BFM).** We initially employ the pixel decoder [46] to gradually upsample visual features  $F_{v_i}$  derived from the SEM to further generate high-resolution per-pixel embeddings  $P_i \in \mathbb{R}^{T \times \frac{H}{2^{i+1}} \times \frac{W}{2^{i+1}} \times C}$ ,  $i \in \{1, 2, 3, 4\}$ , where  $C$  denotes the output channel. Then, we design a Bilateral-Fusion Module (BFM) for constructing a bidirectional audio-visual mapping to assist with segmenting the sounding objects.

As shown in Fig. 3, our BFM utilizes audio features  $F_a \in \mathbb{R}^{T \times D}$ , in conjunction with the largest pixel-level embeddings  $P_1 \in \mathbb{R}^{T \times H_1 \times W_1 \times C}$  as inputs which can propagate ample fine-grained semantic information to audio fea-

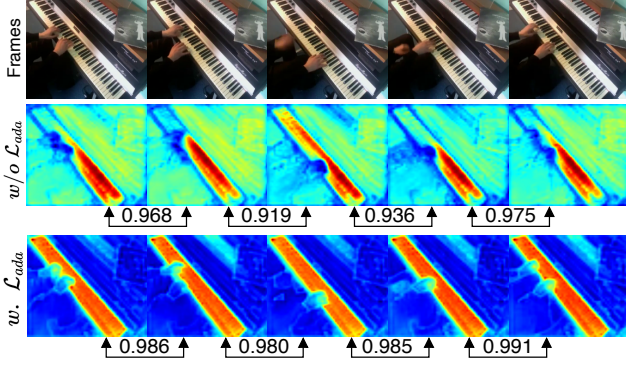


Figure 4. Illustration of the impact on Adaptive Inter-frame Consistency Loss. We visualize the heat map of the predicted masks without and with the consideration of  $\mathcal{L}_{ada}$  based on the S4 subset. The results indicate that implementing  $\mathcal{L}_{ada}$  promotes superior interframe consistency. Best viewed in color.

tures. Here,  $H_1 = H/4$ ;  $W_1 = W/4$ . In order to incorporate both two signals, bilateral attention is designed within our BFM. Specifically, we initially add fixed sine spatial positional encodings and learnable positional encodings to  $P_1$  and  $F_a$ , respectively. Next, in order to integrate audio-visual modalities in a memory-efficient way, our BFM comprises four point-wise linear layers that map  $P_1$  and  $F_a$  to intermediate representations with dimension  $d$ . These representations share queries and keys with queries  $Q = P_1 W_Q$ , keys  $K = F_a W_K$ , visual values  $V_v = P_1 W_V^v$ , and audio values  $V_a = P_1 W_V^a$ . Following the mapping process, the bilateral attention is as follows:

$$P_1 = \text{Softmax}(QK^T/\sqrt{d})V_a + P_1, \quad (4)$$

$$F_a = \text{Softmax}((QK^T/\sqrt{d})^T)V_v + F_a, \quad (5)$$

where  $d$  is the embedding dimension. And  $(QK^T/\sqrt{d})$  is calculated only once, which is more efficient.

After BFM, we proceed by expanding fused audio features  $F_a$  added with learnable embeddings within the transformer decoder as object queries. Additionally, following [10], we generate the output classes  $O^{cls}$  by incorporating the per-pixel embeddings  $P_4, P_3, P_2$  into the transformer decoder. We acquire the predicted masks  $O^{mask}$  by multiplying the output embeddings from the transformer decoder with the fused embedding  $P_1$ .

### 3.3. Mining Temporal Relationships

**Adaptive Inter-frame Consistency Loss.** Temporal always implies a bilateral relationship in nature. For instance, in video clips, one can deduce the scenario of the current frame based on the past frame. Similarly, it is also feasible to predict the future frame based on the current frame. This interactive relations among frames can be construed as a type of temporal entanglement. To take advantage of this

potential temporal entanglement, we introduce an adaptive inter-frame consistency loss for AVS. The similarity score for each successive frame can be calculated as follows:

$$\mathcal{S}_{t:t+1} = \cos(O_t^{mask}, O_{t+1}^{mask}), \quad (6)$$

where  $O_t^{mask}$  refers to the predicted masks at frame  $t$ , and  $\cos(\cdot)$  symbolizes the cosine similarity function. The term  $\mathcal{S}_{t:t+1}$  represents the similarity score between frames  $t$  and  $t+1$ . As illustrated in Fig. 4, it is evident that a significant similarity exists between distinct frames. Therefore, to leverage this prior information, we propose an adaptive inter-frame consistency loss, formulated as follows:

$$\mathcal{L}_{ada} = \sum_{t=1}^{T-1} \exp(\mathcal{S}_{t:t+1} - 1)(1 - \mathcal{S}_{t:t+1}), \quad (7)$$

where  $\exp(\mathcal{S}_{t:t+1} - 1)$  represents adaptive weight. When the disparity between adjacent frames is substantial, the adaptive weight item is minimal, aligning with intuition.

### 3.4. Training and Inference

**Overall Training Loss.** The comprehensive training loss comprises three components: classification loss, mask loss, and adaptive inter-frame consistency loss, as previously discussed. The classification loss is formulated by a cross-entropy loss, denoted as  $\mathcal{L}_{cls} = \mathcal{L}_{ce}$ . The mask loss integrates the binary cross-entropy loss and the dice loss [32], and is depicted as  $\mathcal{L}_{mask} = \mathcal{L}_{bce} + \mathcal{L}_{dice}$ . Considering that in the AVS task, the ratio of segmented objects to the total image area is relatively small, employing dice loss allows the model to better concentrate on the foreground and minimizes distraction from the background. The overall training loss is expressed as follows:

$$\mathcal{L} = \lambda_{cls}\mathcal{L}_{cls} + \lambda_{mask}\mathcal{L}_{mask} + \lambda_{ada}\mathcal{L}_{ada}, \quad (8)$$

where  $\lambda_{cls}$ ,  $\lambda_{mask}$ , and  $\lambda_{ada}$  are hyperparameters. More details about the  $\lambda$  parameters are available in Sec. 4.2.

**Semantic Inference.** After obtaining the predicted embeddings  $O^{cls} \in \mathbb{R}^{T \times N_q \times (K_c+1)}$  and binary masks  $O^{mask} \in \mathbb{R}^{T \times N_q \times H \times W}$ , where  $K_c$  represents the total number of object classes and  $N_q$  is the number of object queries, we employ the identical post-processing as in [10] to yield the final semantic segmentation outputs. Specifically, we first calculate the output mask with classes  $O = O^{cls} \times O^{mask} \in \mathbb{R}^{T \times (K_c+1) \times H \times W}$ . Then, we execute  $\arg \max$  and discard the *no object* class to obtain the ultimate results.

## 4. Experiments

### 4.1. AVSBench Datasets

We evaluate our proposed method on the AVSBench dataset [45], which consists of two scenarios: AVSBench-object and AVSBench-semantic.

Method	Backbone	S4		MS3	
		$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
LVS [5]	ResNet-18	37.9	51.0	29.5	33.0
MSSL [34]	ResNet-18	44.9	66.3	26.1	36.3
3DC [28]	ResNet-152	57.1	75.9	36.9	50.3
SST [12]	ResNet-101	66.3	80.1	42.6	57.2
iGAN [29]	ResNet-50	61.6	77.8	42.9	54.4
LGVT[41]	Swin-B	74.9	87.3	40.7	59.3
AVSBench [45]	ResNet-50	72.8	84.8	47.9	57.8
	PVT-v2	78.7	87.9	54.0	64.5
CSMF [3]	ViT-B	58.0	67.0	34.0	44.0
AVS-BiGen [15]	ResNet-50	74.1	85.4	45.0	56.8
	PVT-v2	81.7	90.4	55.1	66.8
CATR [23]	ResNet-50	74.8	86.6	52.8	65.3
	PVT-v2	81.4	89.6	59.0	70.0
DiffusionAVS [30]	ResNet-50	75.8	86.9	49.8	58.2
	PVT-v2	81.4	90.2	58.2	70.9
ECMVAE [31]	ResNet-50	76.3	86.5	48.7	60.7
	PVT-v2	81.7	90.1	57.8	70.8
BAVS [26]	ResNet-50	78.0	85.3	50.2	62.4
	PVT-v2	82.0	88.6	58.6	65.5
AVSegFormer [13]	ResNet-50	76.5	85.9	49.5	62.8
	PVT-v2	82.1	89.9	58.4	69.3
<b>COMBO (ours)</b>	<b>ResNet-50</b>	<b>81.7</b>	<b>90.1</b>	<b>54.5</b>	<b>66.6</b>
		(+3.7)	(+4.8)	(+2.7)	(+1.3)
	<b>PVT-v2</b>	<b>84.7</b>	<b>91.9</b>	<b>59.2</b>	<b>71.2</b>
		(+2.6)	(+2.0)	(+0.2)	(+1.2)

Table 1. Quantitative comparison results of different methods on AVSBench-object (Single-source, S4; Multi-source, MS3). We use the same backbones (ResNet-50 and PVT-v2) to demonstrate that our method outperforms other methods significantly.

Method	Backbone	AVSS	
		$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
3DC [28]	ResNet-18	17.3	21.6
AOT [39]	ResNet-50	25.4	31.0
AVSBench [45]	ResNet-50	20.2	25.2
	PVT-v2	29.8	35.2
BAVS [26]	ResNet-50	24.7	29.6
	PVT-v2	32.6	36.4
AVSegFormer [13]	ResNet-50	24.9	29.3
	PVT-v2	36.7	42.0
<b>COMBO (ours)</b>	<b>ResNet-50</b>	<b>33.3</b>	<b>37.3</b>
		(+8.4)	(+8.0)
	<b>PVT-v2</b>	<b>42.1</b>	<b>46.1</b>
		(+5.4)	(+4.1)

Table 2. Quantitative comparison results on AVSBench-semantic.

**AVSBench-object.** AVSBench-object [44] is an audio-visual dataset specifically designed for sound target segmentation. AVSBench-object includes two scenarios based on the number of audio sources in each frame: single sound source segmentation (S4) and multiple sound source segmentation (MS3). The S4 scenario incorporates 4,932

Module	S4		AVSS	
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
<b>COMBO</b>	<b>81.7</b>	<b>90.1</b>	<b>33.3</b>	<b>37.3</b>
w/o Siam-Encoder	80.6	88.7	31.9	35.7
w/o Bilateral-Fusion	81.1	89.9	33.1	36.7
w/o Inter-Frame Loss	81.0	89.8	33.0	37.1

Table 3. Ablation study of the various modules included in COMBO. We sequentially remove our proposed modules and compare their performance.

videos, with the ratio of train/validation/test split ratio configured at 70/15/15. This scenario is trained in a semi-supervised manner, wherein each video comprises five frames, but annotation during training is limited to the first frame only. Conversely, the MS3 scenario is characterized by multiple sound sources, including 424 videos, and maintains the same split ratios as in the S4 scenario. This scenario, unlike S4, employs a fully supervised training approach with all five frames being annotated.

**AVSBench-semantic.** AVSBench-semantic [45] is an extension to AVSBench-object that incorporates additional semantic labels for the purpose of enhancing audio-visual semantic segmentation (AVSS). AVSBench-semantic includes a set of new multi-source videos as well as the original AVSBench-object videos, cumulatively accounting for a total of 11,356 videos spanning 70 categories. These videos are allocated as follows: 8,498 for training, 1,304 for validation, and 1,554 for testing. Additionally, the original videos have only been enhanced with semantic information, maintaining the same frames as prior. However, the new videos have been extended to 10 frames, increasing the difficulty due to the inclusion of extended audio-visual sequences.

## 4.2. Experimental Setup

**Implementation Details.** For fair comparison, we adopt the ImageNet pre-trained ResNet-50 [16] and Pyramid Vision Transformer (PVT-v2) [36] as the visual siam-encoders. All input frames are resized to  $224 \times 224$ . In terms of the proposal generator, we leverage the Semantic-SAM [22] to obtain class-agnostic masks. As for audio input, we take the Vggish encoder [17] pre-trained on AudioSet [14] to extract audio features. Following [10], the Multi-Scale Deformable Attention Transformer (MSDeformAttn) is our default pixel decoder. Besides, we adopt the standard transformer decoder, with  $L = 3$  (i.e., a total of 9 layers) and  $N_q = 100$  as the default. The hyperparameters are set as  $\lambda_{cls} = 2$  and  $\lambda_{mask} = 5$ . For the inter-frame consistency loss, we set  $\lambda_{ada} = 10$  for the AVSBench-object dataset,  $\lambda_{ada} = 5$  for the AVSBench-semantic dataset due to longer frames per video. We calculate the inter-frame consistency loss only in the intermediate transformer decoder layer to prevent an over-dependence on temporal information. All models are trained using the Adam optimizer

SEM Module	S4		AVSS		#Params
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	
<i>w.</i> shared weights	81.5	90.2	31.1	34.4	135.5
<i>w/o</i> shared weights	<b>81.7</b>	<b>90.1</b>	<b>33.3</b>	<b>37.3</b>	158.9

Table 4. Ablation study of Siam-Encoder Module (SEM). We explore two configurations: using shared weights and using separate weights. #Params denotes the model parameters (M).

Fusion Mode	S4		AVSS	
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
<i>w/o</i> fusion	81.1	89.9	33.1	36.7
visual fusion only	81.3	89.6	32.6	36.4
audio fusion only	81.3	<b>90.2</b>	33.2	36.7
<b>bilateral fusion</b>	<b>81.7</b>	90.1	<b>33.3</b>	<b>37.3</b>

Table 5. Ablation study of fusion mode. We conduct among four modes: no fusion, visual fusion only, audio fusion only, and our fusion mode.

with a learning rate of  $1e-4$  and weight decay of 0.05. We train the S4 and AVSS subsets for 90k and MS3 for 20k iterations with a batch size of 8 on a single A100 40GB GPU. **Metrics.** Following [44], we adopt two metrics to verify the effectiveness, namely, Jaccard index  $\mathcal{J}$  and F-score  $\mathcal{F}$ .  $\mathcal{J}$  computes the intersection over union (IoU) of the predicted segmentation and the ground truth mask.  $\mathcal{F}$  considers both precision and recall, which is represented as  $\mathcal{F}_{\beta} = \frac{(1+\beta^2) \times \text{precision} \times \text{recall}}{\beta^2 \times \text{precision} + \text{recall}}$ , where  $\beta^2$  is set at 0.3. In our experiment, we use  $\mathcal{M}_{\mathcal{J}}$  and  $\mathcal{M}_{\mathcal{F}}$  to denote the mean metrics across the entire dataset.

### 4.3. Main Results

We conduct experiments on AVSBench-object and AVSBench-semantic datasets. As AVS is an emerging proposed problem recently introduced by [44], we compare our COMBO with some state-of-the-art methods from other related tasks, such as sound source localization (SSL) [5, 34], video object segmentation (VOS) [12, 28, 39], and salient object detection (SOD) [29, 41], all of which provide a comparative benchmark for our experiments. As evidenced in Tab. 1, COMBO demonstrates a substantial performance gap (+9.8 mIoU in S4; +16.3 mIoU in MS3) over other related methods, principally attributable to variances in setting specific task scenarios. We also compare our method against some recent state-of-the-art methods [13, 15, 23, 26, 30, 31, 44] that have been explicitly designed for audio-visual segmentation settings. On the AVSBench-object dataset, COMBO-R50 outperforms the current best performance by achieving 3.7 mIoU and 4.8 F-score improvements for S4 subset and 2.7 mIoU and 1.3 F-score improvements for MS3, while COMBO-PVT surpasses the top-performing method by 2.6 mIoU and 2.0 F-score for S4 and 0.2 mIoU and 1.2 F-score for MS3.

Besides, we compare the AVSBench-semantic dataset as

Queries	S4		AVSS	
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
all	80.9	89.9	30.7	34.2
<b>add</b>	<b>81.7</b>	<b>90.1</b>	<b>33.3</b>	<b>37.3</b>

Table 6. Ablation study of learnable queries.

$\lambda_{ada}$	S4		AVSS	
	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$	$\mathcal{M}_{\mathcal{J}}$	$\mathcal{M}_{\mathcal{F}}$
0	81.0	89.8	33.0	37.1
5	81.1	<b>90.1</b>	<b>33.3</b>	<b>37.3</b>
10	<b>81.7</b>	<b>90.1</b>	32.1	35.6
20	81.2	89.7	32.6	36.1

Table 7. Ablation study of the adaptive inter-frame consistency loss.  $\lambda_{ada}$  is the hyperparameter, while higher values constrain the output to be more similar.

displayed in Tab. 2, which presents a more challenging setting. Both COMBO-R50 and COMBO-PVT achieve significant results, with 8.4 and 5.4 mIoU improvements and significant F-score enhancements of 8.0 and 4.1, respectively. These experiments confirm that our COMBO model surpasses existing state-of-the-art methods across all sub-tasks, consequently setting a new benchmark for AVS.

### 4.4. Ablation Study

In this section, we conduct ablation studies to verify the effectiveness of each essential design in the proposed COMBO. Specifically, we adopt ResNet-50 [16] as the backbone and carry out extensive experiments on the S4 and AVSS sub-tasks due to more videos in these tasks. Other training settings remain consistent with Sec. 4.2

**Component analysis of COMBO.** To validate the impact of our proposed method, we separately eliminate the Siam-Encoder Module (SEM), Bilateral-Fusion Module (BFM), and adaptive inter-frame consistency loss. As demonstrated in Tab. 3, the results indicate that COMBO has demonstrated superior influence on SEM with 1.4 mIoU improvement, particularly on the AVSS subset. Concurrently, our findings indicate that BFM is of substantial significance, demonstrating a performance enhancement of 0.6 mIoU over the S4 subset. More analysis are discussed later. In addition, we examine the effect of inter-frame consistency loss. The results reveal that our loss function contributes to performance improvements with 0.7 mIoU on the S4 subset. More details are provided subsequently.

**Effects of Siam-Encoder Module (SEM).** We first examine the significance of our SEM. Two designs are compared in Tab. 4: one with share weights and another with separate weights. When comparing shared and separate parameters, it is found that separable parameters can attain 0.2 mIoU on the S4 subset, but it is indeed more costly than the others.

**Effects of Bilateral-Fusion Module (BFM).** Moreover, we investigate the influence of our BFM. We compared our

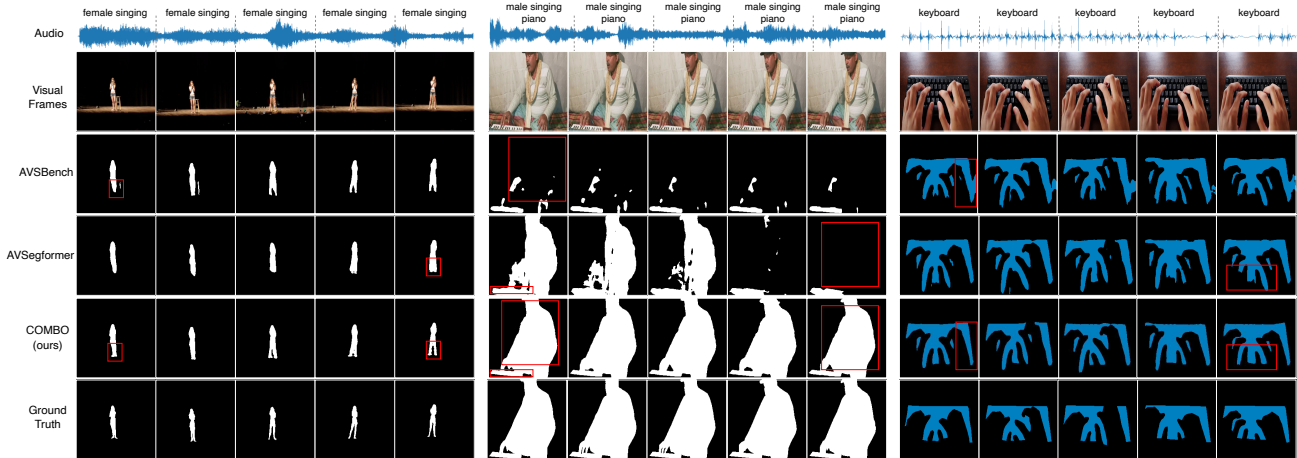


Figure 5. Comparison of Visual Examples on the AVSBench-object and AVSBench-semantic Datasets with AVSBench [45] and AVSegformer [13]. Wherein the leftmost example is derived from the S4 subset, the middle example is from the MS3 subset, and the rightmost example is from the AVSS subset. Red bounding boxes highlight the specific regions for comparison.

modules using different variants, which include without any fusion, video-only fusion (inject audio feature into visual), audio-only fusion (inject visual feature to audio), and our bilateral fusion. As demonstrated in Tab. 5, bilateral fusion achieves a performance improvement of 0.6 mIoU compared to the model without any fusion. It is worth noting that the transformer decoder inherently consists of a cross-attention function, potentially serving as a fusion process. Our model also realizes a performance improvement of 0.4 mIoU compared to only audio or visual fusion.

**Effects of Audio Queries.** As demonstrated in Tab. 6, *All* denotes the exclusive use of fused audio queries, and *Add* signifies the combination of fused audio queries and learnable queries. We first expand the audio features to the exact dimensions as learnable queries, then compare the experiments with the exclusive use of fused audio queries and the combination of fused audio queries and learnable queries. The results show that queries with *Add* have 0.8 and 2.6 mIoU improvements over *All* on S4 and AVSS subsets.

**Effects of Adaptive Inter-frame Consistency Loss.** To validate the effect of  $\mathcal{L}_{ada}$ , we adopt different values of  $\lambda_{ada}$ . As shown in Tab. 7, the experiments demonstrate that appropriate consistency constraints can enhance the model’s performance. The appropriate value  $\lambda_{ada}$  can improve performance by 0.7 mIoU and 0.3 F-score for S4. In addition, the heat map of the predicted masks has been visualized in Fig. 4. It is observable that the use of  $\mathcal{L}_{ada}$  facilitates the generation of a more distinct boundary output. Nevertheless, it is essential to note that exceedingly high values may lead to a decline in performance, given that the video does not exhibit complete consistency.

#### 4.5. Qualitative Analysis

We provide a qualitative comparison between AVSBench [45], AVSegformer [13] and our proposed method

on AVSBench-object and AVSBench-semantic datasets. As depicted in Fig. 5, our method, COMBO, exhibits superior audio-temporal and spatial localization quality, leading to better visualization and segmentation performance. For instance, in the case of the middle samples, our model accurately segments the singing man despite the presence of other sounds. Moreover, our method achieves more precise segmentation for background noise handling and provides richer details of the foreground in other examples.

## 5. Conclusion

We introduce a novel audio-visual transformer framework, termed COMBO, that archives state-of-the-art performance on AVSBench-object and AVSBench-semantic datasets. Contrary to previous methodologies that only factor in modality or temporal relations individually, our method explores multi-order bilateral relations for the first time, combining pixel entanglement, modality entanglement, and temporal entanglement. For these three kinds of entanglement, we propose Siam-Encoder Module (SEM), Bilateral-Fusion Module (BFM), and adaptive inter-frame consistency loss, respectively. Extensive experimental results verify the effectiveness of our proposed framework. We hope that our work will inspire further research in this significant and worthwhile field.

## Acknowledgements

This research was supported by the Strategic Priority Research Program of Chinese Academy of Sciences (Grant No. XDB0500103), the National Natural Science Foundations of China (Grants No. 62076242, 62376267), the Pre-Research Project on Civil Aerospace Technologies (No. D030312), the National Defense Basic Scientific Research Program of China (No. JCKY2021203B063) and the innoHK project.



## References

- [1] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE international conference on computer vision*, pages 609–617, 2017. 1
- [2] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. *Advances in neural information processing systems*, 29, 2016. 1
- [3] Swapnil Bhosale, Haosen Yang, Diptesh Kanojia, and Xiatian Zhu. Leveraging foundation models for unsupervised audio-visual segmentation. *arXiv preprint arXiv:2309.06728*, 2023. 2, 3, 6
- [4] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 721–725. IEEE, 2020. 1
- [5] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Localizing visual sounds the hard way. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16867–16876, 2021. 1, 2, 6, 7
- [6] Jiaqi Chen, Jiachen Lu, Xiatian Zhu, and Li Zhang. Generative semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7111–7120, 2023. 4
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2018. 2
- [8] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. arxiv 2017. *arXiv preprint arXiv:1706.05587*, 2, 2019. 2
- [9] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 2
- [10] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1280–1289, 2022. 5, 6
- [11] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 2
- [12] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W. Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5912–5921, 2021. 6, 7
- [13] Shengyi Gao, Zhe Chen, Guo Chen, Wenhai Wang, and Tong Lu. Avsegformer: Audio-visual segmentation with transformer. *arXiv preprint arXiv:2307.01146*, 2023. 2, 3, 6, 7, 8
- [14] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 776–780, 2017. 6
- [15] Dawei Hao, Yuxin Mao, Bowen He, Xiaodong Han, Yuchao Dai, and Yiran Zhong. Improving audio-visual segmentation with bidirectional generation. *arXiv preprint arXiv:2308.08288*, 2023. 2, 3, 6, 7
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 6, 7
- [17] Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, R. Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron J. Weiss, and Kevin Wilson. Cnn architectures for large-scale audio classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 131–135, 2017. 4, 6
- [18] Di Hu, Rui Qian, Minyue Jiang, Xiao Tan, Shilei Wen, Errui Ding, Weiyao Lin, and Dejing Dou. Discriminative sounding objects localization via self-supervised audiovisual matching. *Advances in Neural Information Processing Systems*, 33:10077–10087, 2020. 1, 2
- [19] Jitesh Jain, Jiachen Li, Mang Tik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. Oneformer: One transformer to rule universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2989–2998, 2023. 2
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4026, 2023. 2, 3
- [21] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. *arXiv preprint arXiv:2308.00692*, 2023. 3
- [22] Feng Li, Hao Zhang, Peize Sun, Xueyan Zou, Shilong Liu, Jianwei Yang, Chunyuan Li, Lei Zhang, and Jianfeng Gao. Semantic-sam: Segment and recognize anything at any granularity. *arXiv preprint arXiv:2307.04767*, 2023. 2, 3, 6
- [23] Kexin Li, Zongxin Yang, Lei Chen, Yi Yang, and Jun Xun. Catr: Combinatorial-dependence audio-queried transformer for audio-visual video segmentation. *arXiv preprint arXiv:2309.09709*, 2023. 2, 3, 6, 7
- [24] Liumian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. Grounded language-image pre-training. In

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10965–10975, 2022. 2
- [25] Yuhang Ling, Yuxi Li, Zhenye Gan, Jiangning Zhang, Mingmin Chi, and Yabiao Wang. Hear to segment: Unmixing the audio to guide the semantic segmentation, 2023. 2
- [26] Chen Liu, Peike Li, Hu Zhang, Lincheng Li, Zi Huang, Dadong Wang, and Xin Yu. Bavs: Bootstrapping audio-visual segmentation by integrating foundation knowledge. *arXiv preprint arXiv:2308.10175*, 2023. 2, 3, 6, 7
- [27] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015. 2
- [28] Sabarinath Mahadevan, Ali Athar, Aljosa Osep, Laura Leal-Taixé, Bastian Leibe, and Sebastian Hennen. Making a case for 3d convolutions for object segmentation in videos. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020. 6, 7
- [29] Yuxin Mao, Jing Zhang, Zhexiong Wan, Yuchao Dai, Aixuan Li, Yunqiu Lv, Xinyu Tian, Deng-Ping Fan, and Nick Barnes. Generative transformer for accurate and reliable salient object detection. *arXiv preprint arXiv:2104.10127*, 2021. 6, 7
- [30] Yuxin Mao, Jing Zhang, Mochu Xiang, Yunqiu Lv, Yiran Zhong, and Yuchao Dai. Contrastive conditional latent diffusion for audio-visual segmentation. *arXiv preprint arXiv:2307.16579*, 2023. 3, 6, 7
- [31] Yuxin Mao, Jing Zhang, Mochu Xiang, Yiran Zhong, and Yuchao Dai. Multimodal variational auto-encoder based audio-visual segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 954–965, 2023. 2, 3, 6, 7
- [32] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571, 2016. 5
- [33] Shentong Mo and Yapeng Tian. Av-sam: Segment anything model meets audio-visual localization and segmentation. *arXiv preprint arXiv:2305.01836*, 2023. 2, 3
- [34] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. In *Computer Vision – ECCV 2020*, pages 292–308, Cham, 2020. Springer International Publishing. 1, 2, 6, 7
- [35] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7242–7252, 2021. 2
- [36] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):415–424, 2022. 6
- [37] Yaoting Wang, Weisong Liu, Guanyao Li, Jian Ding, Di Hu, and Xi Li. Prompting segmentation with sound is generalizable audio-visual source localizer. *arXiv preprint arXiv:2309.07929*, 2023. 2, 3
- [38] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. 2, 4
- [39] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34:2491–2502, 2021. 6, 7
- [40] 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. 2, 3
- [41] Jing Zhang, Jianwen Xie, Nick Barnes, and Ping Li. Learning generative vision transformer with energy-based latent space for saliency prediction. In *Advances in Neural Information Processing Systems*, pages 15448–15463. Curran Associates, Inc., 2021. 6, 7
- [42] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6230–6239, 2017. 2
- [43] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H.S. Torr, and Li Zhang. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6877–6886, 2021. 2
- [44] Jinxing Zhou, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, and Yiran Zhong. Audio-visual segmentation. In *European Conference on Computer Vision*, pages 386–403. Springer, 2022. 1, 2, 3, 4, 6, 7
- [45] Jinxing Zhou, Xuyang Shen, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, et al. Audio-visual segmentation with semantics. *arXiv preprint arXiv:2301.13190*, 2023. 1, 3, 5, 6, 8
- [46] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable {detr}: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations*, 2021. 4