

Towards Robust Audiovisual Segmentation in Complex Environments with Quantization-based Semantic Decomposition

Xiang Li¹, Jinglu Wang², Xiaohao Xu³, Xiulian Peng², Rita Singh¹, Yan Lu², Bhiksha Raj¹
¹ CMU, ² Microsoft, ³ UMich

Abstract

Audiovisual segmentation (AVS) is a challenging task that aims to segment visual objects in videos according to their associated acoustic cues. With multiple sound sources and background disturbances involved, establishing robust correspondences between audio and visual contents poses unique challenges due to (1) complex entanglement across sound sources and (2) frequent changes in the occurrence of distinct sound events. Assuming sound events occur independently, the multi-source semantic space can be represented as the Cartesian product of single-source subspaces. We are motivated to decompose the multi-source audio semantics into single-source semantics for more effective interactions with visual content. We propose a semantic decomposition method based on product quantization, where the multi-source semantics can be decomposed and represented by several disentangled and noise-suppressed single-source semantics. Furthermore, we introduce a global-to-local quantization mechanism, which distills knowledge from stable global (clip-level) features into local (frame-level) ones, to handle frequent changes in audio semantics. Extensive experiments demonstrate that our semantically decomposed audio representation significantly improves AVS performance, e.g., +21.2% mIoU on the challenging AVS-Semantic benchmark with ResNet50 backbone. <https://github.com/lxa9867/QSD>.

1. Introduction

Prompt-guided video object segmentation (VOS) [18, 31, 61, 63, 68, 69] focuses on the segmentation of objects in videos according to prompts and predict the object classes if required. The use of various modalities of prompts, such as *text* and *audio*, has promising implications for multi-modal understanding and user-interactive video applications.

Numerous recent referring VOS methods [7, 61–63] have demonstrated their effectiveness in using linguistic prompts that describe objects discriminatively at a video clip level. However, employing audio prompts for active-sounding objects in audiovisual segmentation (AVS) [68,

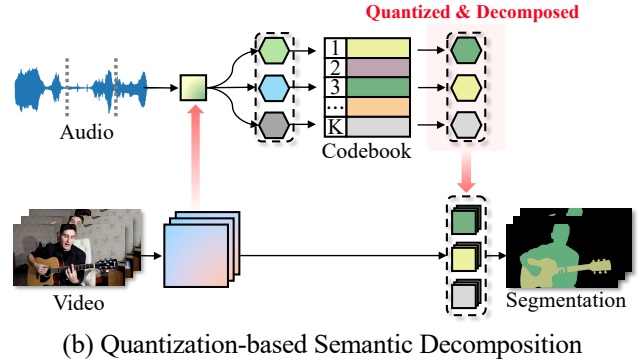
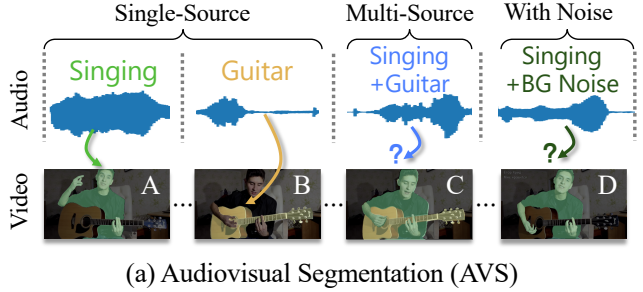


Figure 1. (a) AVS aims to segment masks associated with sound-ing objects at each timestamp and predict their classes if required. It presents substantial challenges with multi-source audio (Frame C) or background disturbances (Frame D), due to feature entangle-ment and noise. (Provided class labels are for visualization only and not included in the input.) (b) Our method tackles challenges by decomposing audio features into several semantic tokens. We employ product quantization (with a shared codebook for each de-composed feature) as an information bottleneck to achieve a com-pact and informative representation.

69] introduces unique challenges. First, multiple audio sources could **entangle at any timestamp**, meaning that audio in a single frame could correspond to multiple segmen-tation masks of different classes, as shown in Frame C of Fig. 1 (a). This creates ambiguity in aligning mixed-source audio with visual content, adding complexity to the AVS problem. Second, audio is often accompanied by **back-ground disturbances**, such as noise or sounds from ob-




Figure 2. **An example of decomposing a two-source semantic space.** A multi-source audio semantic space can be assumed as a Cartesian product of two identical single-source semantic spaces, which can be decomposed via product quantization. The red points represent the quantized semantics.

jects outside the frame, as shown in Frame D of Fig. 1 (a). These disturbances can intensify the difficulties in AVS since acoustic cues are inherently susceptible to such interference. As such, directly interacting visual content with audio in complex environments (noisy and entangled) [21, 34, 68] by imitating clean and discriminative text prompts like [7, 62] is not effective. Moreover, adopting pre-trained models for audio separation and noise suppression in advance requires additional data, lacks the flexibility for fine-tuning on AVS-specific datasets, and fails to consider the impact of related visual content. Addressing these issues by seeking a compact and source-disentangled representation for audio could greatly improve audiovisual interaction and further simplify the AVS problem. Also, we believe that exploring solutions to the challenges posed by multi-source and noisy prompts could contribute to the broader field of multi-modal understanding.

In this paper, we propose a novel AVS method that leverages the information bottleneck (IB) principle [1, 3, 58] to learn a compact and semantically decomposed audio representation, and then query them to visual features to produce masks, as illustrated in Fig. 1 (b). The IB principle is beneficial as it seeks a compact and task-relevant representation that discards unnecessary details, aids in disentangling underlying factors, and simplifies matching and comparison by converting complex features into discrete symbols. Specifically, we employ a Product Quantization (PQ) based method to decompose multi-source semantics thanks to its ability to represent a complex space through the product of several subspaces [30]. We define the *audio semantic space* associated with the label set of audio at all timestamps, and the *visual semantic space* associated with the label set of all pixels. An example of decomposing semantic space is illustrated in Fig. 2. During single-source moments (Frame A/B of Fig. 1 (a)), the visual semantic space corresponds to the semantic labels $\mathcal{Y} = \{\text{Guitar (G), Singing (S)}\}$, and the audio semantic space

of the same label set \mathcal{Y} . During the two-source moment (Frame C of Fig. 1 (a)), the visual semantic space remains consistent to \mathcal{Y} , but the *possible* two-source audio semantic space expands to the label set $\mathcal{Y} \times \mathcal{Y} = \{\text{GG, GS, SG, SS}\}$. We simplify the AVS problem by assuming independent sound events, which allows us to represent the multi-source semantic space as a Cartesian product of identical single-source semantic spaces. By applying product quantization as an information bottleneck and constraining representation dimension to be much smaller than the size of semantic space, achieved by utilizing a shared codebook, we can decompose the multi-source audio semantics into single-source semantics with noise suppression.

Furthermore, it is necessary to analyze audio semantics for every frame, given that sounding sources can change frequently (as the example in Fig. 1 (a)). This presents a robustness challenge, as extracted semantics from short-term audio can be unstable compared to the extraction from long-term (clip-level) audio. To improve the frame-level audio representation, we propose a global-to-local mechanism, which distills knowledge from robust global (clip-level) audio representations into local (frame-level) ones. Specifically, the codebook is learned with clip-level visual-enriched audio features and performs local quantization on each frame without updating it. Thereby, the local semantic tokens are calibrated to the more robust and representative clip-level feature.

In summary, our contribution is three-fold:

- An effective approach of audio semantic decomposition via product quantization, addressing the challenge of correlating audiovisual features introduced by multiple sounding sources and background disturbances.
- A global-to-local distilling mechanism for frame-level audio semantic enhancement, addressing the ineffectiveness of frame-level audio feature extraction.
- Extensive experiments are conducted to verify the effectiveness and robustness of the proposed method, which significantly outperforms previous state-of-the-art methods on three AVS benchmarks (+21.2% mIoU for AVS-Semantic with ResNet-50 backbone).

2. Related Work

Audiovisual segmentation and localization. Audiovisual segmentation (AVS), which was recently introduced [68], aims to segment the objects that produce sound at the time of the image frame. Zhou *et al.* [68] proposed a method with cross-modal attention to locate the sound source, making it the pioneering work in AVS. Recently, an extended task of AVS, audiovisual semantic segmentation (AVSS), is proposed by Zhou *et al.* [69] which aims to not only segment the mask of sound sources but also predict the category of each sound source. Due to the semantic entanglement in audio, tackling multi-source AVSS is more chal-

lenging than AVS task. Zhou *et al.* [69] follows the TPAVI module in [68] to conduct audiovisual interaction. More recently, CATR [34] introduces an encoding-decoding framework CATR that presents a novel spatial-temporal audio-video fusion block to fully consider the audio-visual combinatorial dependence in a decoupled and memory-efficient manner. Chen *et al.* [66] creates a novel benchmark called VPO which augments images from COCO [38] with VGGSound [8] to facilitate the AVS research. Sound source localization (SSL) [2, 9, 26, 45, 46, 50, 52] is a related problem to AVS that aims to locate the regions of sounds in the visual frame. Common SSL methods [4, 5, 17, 52] leverage cross-modal correspondence between audio and visual features to locate sounds, which are then displayed as heatmaps. For instance, Mo *et al.* [45] leverage multi-level audiovisual contrastive learning to effectively locate the objects. Different from previous methods primarily designed for single-source scenarios, our objective is to address the semantic entanglement present in multi-source audios and explore methods for effective interaction between multi-source audios and videos.

Video Segmentation. Audiovisual segmentation is intricately connected to video object segmentation (VOS) [15, 16, 27, 29, 54, 55, 64, 65] and video semantic segmentation (VSS) [28, 37, 48, 56, 70]. Notably, a most closely related task in the video segmentation domain is referring video object segmentation (R-VOS) [6, 12, 23, 33, 36, 44, 53, 57, 60–62], which seeks to segment objects in visual frames based on linguistic expressions. In R-VOS, each expression typically refers to a single object, akin to the single-source AVS setting. Recently, Ding *et al.* [19] introduced the MeViS dataset, expanding the R-VOS task to a more generalized setting that allows referring to multiple objects through a single language expression.

3. Method

In this section, we first present the formulation of the quantization-based method for multi-source audio semantic decomposition. Then, we outline the pipeline that utilizes the quantized and decomposed audio representation to improve the audiovisual segmentation tasks.

3.1. Quantization-based Semantic Decomposition

Without loss of generality, let us consider the N -source audio case at a specific timestamp, where each source corresponds to a label in the single-source label set \mathcal{Y}_s . Due to the entanglement, the multi-source audio feature $x \in \mathcal{X}_m$ will correspond to a much larger label set \mathcal{Y}_m , where \mathcal{X}_m denotes the multi-source audio feature space.

Intuitively, x containing information from N -source audio is highly entangled and not effective for interacting with visual content. Our objective is to seek a decomposition

transformation $\pi(x) = (x_1, \dots, x_N)$. We expect that the decomposed N representations (x_1, \dots, x_N) can be disentangled, such that each x_i can retain the most relevant information to the single-source label in \mathcal{Y}_s while suppressing extraneous and entangled noise. Inspired by the information bottleneck principle, which facilitates data compression to encapsulate the most task-relevant information and break down the data into independent factors, we propose to apply a quantization-based semantic decomposition method.

Assuming the independence of sound events, \mathcal{Y}_m becomes a Cartesian product of several equivalent single-source label sets \mathcal{Y}_s . We have:

$$\mathcal{Y}_m = \underbrace{\mathcal{Y}_s \times \dots \times \mathcal{Y}_s}_N. \quad (1)$$

Knowing the desired x_i should effectively represent the semantics within \mathcal{Y}_s , we construct an information bottleneck by constraining each x_i for $i = 1, \dots, N$, to exist within the same space, a space with cardinality K roughly equivalent to the size of \mathcal{Y}_s . Specifically, we perform the product quantization [30] on $x \in \mathcal{X}_m$ with N shared vector quantizer $VQ(\cdot)$ as:

$$PQ(x) = VQ(x_1) \oplus \dots \oplus VQ(x_N), \quad VQ \sim \mathcal{C}, \quad (2)$$

where \oplus denotes channel-wise concatenation for recombining the quantized features, \sim denotes that the vector quantizer $VQ(\cdot)$ is associated with the codebook $\mathcal{C} = \{e^k\}_{k=1}^K$, that is, $VQ(\cdot)$ maps a feature x_i to a codeword $e_i = \arg \min_{e^k \in \mathcal{C}} \|x_i - e^k\|_p$ that minimizes the distance between x_i and $e^k \in \mathcal{C}$ in the p -norm sense. In this way, with task supervision to enforce $PQ(x)$ to represent the same information as x , we can effectively obtain decomposed x_i using the constructed information bottleneck.

3.2. Network Overview

The proposed framework consists of three main components as illustrated in Fig. 3: feature encoding, global decomposition, and local calibration.

First, we extract visual features $F_v = \{f_{v,t}\}_{t=1}^T$ with a visual encoder (a backbone with a transformer encoder on top of the backbone), and acoustic features $F_a = \{f_{a,t}\}_{t=1}^T$ with VGGish [25] (the same setting as previous AVS methods [34, 68]). Then, to decompose semantics in multi-source audio features, we use a global semantic decomposition module to map the audio query into a set of semantic tokens $\{g_i\}_{i=1}^N$. We learn a semantic codebook to quantize them. The quantized tokens are further employed to modulate the visual features to inject information about corresponding sound sources. Afterward, to obtain frame-level audio features to query object masks, we utilize a local semantic decomposition module for each timestamp, which uses the global codebook to decouple local audio semantics. Each quantized local semantic token $VQ(l_{i,t})$ serves




Figure 3. **Method overview.** Our pipeline accepts an input video along with its associated and non-separated audio, proceeds to segment masks for each frame based on the active audio, and predicts semantic labels of masks if required. The **global decomposition** module decomposes multi-source audio features F'_a and quantizes all decomposed single-source audio features $\{g_i\}$ with a shared codebook. Visual features F_v are further fused with decomposed audio features to be F'_v . The **local calibration** module distills knowledge from stable global (clip-level) audio features $\{VQ(g_i)\}$ to local (frame-level) audio features $\{VQ(l_{i,t})\}$ by utilizing the shared codebook. Masks for each frame are generated by the semantic-guided mask decoder with local audio features $\{VQ(l_{i,t})\}$ and the fused visual feature F'_v .




Figure 4. (a) Global and local audio semantic decoder (ASD) share similar structures that query clip-/frame-level audio features, F'_a or $f_{a,t}$, with a transformer decoder TrD_{global}/TrD_{local} using learnable semantic prototypes $\{p_i\}$. (b) The semantic-guided mask decoder contains a transformer decoder TrD_{segm} to align audiovisual features and computes dynamic filters $\theta_{i,t}$. The final mask $M_{i,t}$ is generated by a dynamic convolution between the visual feature $f_{v,t}^{out}$ and $\theta_{i,t}$.

as a query to segment a frame-level mask with the semantic-guided mask decoder.

3.3. Global Decomposition

To tackle the mixture of multi-source audio queries and effectively conduct audiovisual fusion, we propose the global decomposition module to decompose audio semantics at the video clip level, illustrated in the pink box of Fig. 3. It consists of two stages: global semantic decomposition (shown in the lower part) and audiovisual semantic recombination (shown in the upper part).

Global semantic decomposition. Global semantic decomposition aims to decompose multi-source audio semantics into single-source semantics. The audio feature F_a is first

fused with video feature F_v to be F'_a , taking the form:

$$\begin{aligned} F'_a &= \text{LN}(\text{FFN}(h_a) + h_a), \\ h_a &= \text{LN}(\text{MCA}(F_a, F_v) + F_a), \end{aligned} \quad (3)$$

where MCA denotes Multi-head Cross-Attention, LN denotes Layer Normalization, and FFN denotes Feed-Forward Network. After that, we transform the audio feature F'_a to N decomposed semantic tokens $\{g_i\}_{i=1}^N$ with a global audio semantic decoder (global ASD),

$$g_i = \text{TrD}_{global}(p_i | F'_a), \quad (4)$$

by querying a set of learnable semantic prototypes $\{p_i\}_{i=1}^N$ to the feature F'_a , with a transformer decoder TrD_{global} as shown in Fig. 4 (a). Each semantic token is then quantized to be $e_i = VQ(g_i)$ with the shared codebook $\mathcal{C} = \{e^k\}_{k=1}^K$, imposing that all semantic tokens to share an identical feature subspace with low cardinality. Note that we set the codebook size $K \approx |\mathcal{Y}|$ to force the network to learn decomposed semantics, where $|\mathcal{Y}|$ denotes the number of sound event categories.

Audiovisual semantic recombination. Audiovisual semantic recombination aims to leverage the decomposed audio feature to interact with visual features. After obtaining quantized global semantic tokens $\{e_i\}_{i=1}^N$, which encode N groups of decomposed semantics, we aim to interact them with visual features while preserving the original function of the multi-source audio input. A set of dynamic filters $\{w_i \in \mathbb{R}^{D_v}\}_{i=1}^N$ are first learned from global semantic tokens $\{g_i\}_{i=1}^N$ by two linear layers. D_v is the channel dimension of F_v . After that, we utilize channel-wise attention to modulate video features by each filter to interact the visual

feature with the content referred by different semantic tokens, which is given by:

$$F'_v = \text{BN}(\varphi(w_i F_v \oplus \dots \oplus w_N F_v) + F_v), \quad (5)$$

where φ denotes a convolution layer to reduce channel from $N \times D_v$ to D_v , BN denotes Batch Normalization, and \oplus denotes concatenation among channels. By incorporating channel-wise attention, the visual features can be more effectively concentrated on the relevant audio content. Furthermore, through channel-wise concatenation, the decomposed audio semantics can be reintegrated, producing hybrid semantics that refer to the holistic contents of the original audio input.

3.4. Local Calibration

Since the audio query is time-variant, global semantic tokens cannot be accurately aligned with visual features at the frame level. To segment audio-queried contents in each frame, we propose the local calibration module (illustrated in the blue box of Fig. 3), consisting of a local semantic decomposition stage (lower part) and a semantic-guided mask decoding stage (upper part).

Local semantic decomposition. This stage aims to decompose the semantics encoded in each audio frame. Similar to the global ASD, the local audio semantic decoder (Local ASD) decodes frame-level semantics with a transformer decoder $\text{TrD}_{\text{local}}$ and a set of semantic prototypes $\{p_i\}_{i=1}^N$, as shown in Fig. 4 (a). The local semantic tokens $l_{i,t}$ are given by

$$l_{i,t} = \text{TrD}_{\text{local}}(p_i | f_{a,t}). \quad (6)$$

where $f_{a,t}$ is t -th frame of the audio feature F_a . The local semantic tokens do not build their own codebook but utilize the global codebook \mathcal{C} , that is, they do not update \mathcal{C} but are committed to being close to the vectors in \mathcal{C} . In this way, the local semantic tokens can be calibrated according to global ones, which are more reliable. Further explanation regarding supervision will be provided in Sec. 3.5.

Semantic-guided mask decoding. We utilize the semantic-guided mask decoder to decode visual features into masks that correspond to decomposed local audio semantics, with the detailed structure illustrated in Fig. 4 (b). Pyramid video features $F_v^{\text{out}} = \{f_{v,t}^{\text{out}}\}_{t=1}^T$ are obtained with the feature pyramid network [39]. We leverage a shared multimodal transformer decoder TrD_{segm} to generate dynamic filters $\theta_{i,t} = \phi_{\text{segm}}(\text{TrD}_{\text{segm}}(\text{VQ}(l_{i,t}) | f_{v,t}))$ for each timestep, where ϕ_{segm} is a two-layer fully-connected network. The final mask segmentation can be obtained by:

$$M_{i,t} = f_t^{\text{out}} * \theta_{i,t}, \quad (7)$$

where $*$ denotes the dynamic convolution [13]. Each filter represents semantics of a decomposed single-source audio,

contributing to the segmentation of the single sounding object. Additional class probability prediction $P_{i,t}$ and bounding box prediction $B_{i,t}$ for each mask $M_{i,t}$ are performed by two two-layer fully connected networks from the output of $\text{TrD}_{\text{segm}}(\text{VQ}(l_{i,t}) | f_t)$.

3.5. Loss Function

The overall loss function is given by

$$\mathcal{L} = \lambda_{\text{quant}} \mathcal{L}_{\text{quant}} + \mathcal{L}_{\text{segm}}, \quad (8)$$

where $\mathcal{L}_{\text{quant}}$ and $\mathcal{L}_{\text{segm}}$ are the loss for semantic quantization and segmentation, respectively. λ_{quant} is a constant.

Loss for semantic quantization. The quantizer is shared with both global and local semantic decomposition, while the local semantic tokens do not update the codebook. The loss is given by

$$\begin{aligned} \mathcal{L}_{\text{quant}} = & \sum_{i=1}^N \left\{ \|\text{VQ}(g_i) - \text{sg}[g_i]\|_2^2 + \lambda_{\text{com}} \|\text{sg}[\text{VQ}(g_i)] - g_i\|_2^2 \right. \\ & \left. + \lambda_{\text{com}} \sum_{t=1}^T \|\text{sg}[\text{VQ}(l_{i,t})] - l_{i,t}\|_2^2 \right\}, \end{aligned} \quad (9)$$

where $\text{sg}[\cdot]$ stands for stop-gradient operation. $\text{VQ}(\cdot)$ denotes the vector quantization function, where $\text{VQ}(x) = e_i = \arg \min_{e^k} \|x - e^k\|_2 \in \mathcal{C}$ and $\mathcal{C} = \{e^k\}_{k=1}^K$ is the shared codebook. The first term aims to update the codebook. The second and third terms aim to minimize the quantization error by forcing the input vector to be quantized to its closest vector in the codebook.

Loss for segmentation. Let the predictions of the network be $\mathbf{y} = \{\mathbf{y}_i\}_{i=1}^N$ where $\mathbf{y}_i = \{B_{i,t}, P_{i,t}, M_{i,t}\}_{t=1}^T$. $B_{i,t}$, $P_{i,t}$ and $M_{i,t}$ denote bounding box, class probability and mask predictions, respectively. We denote the ground-truth as $\hat{\mathbf{y}} = \{\hat{\mathbf{y}}_j\}_{j=1}^N$ (padded with \emptyset [14]) where $\hat{\mathbf{y}}_j = \{\hat{B}_{j,t}, \hat{C}_{j,t}, \hat{M}_{j,t}\}_{t=1}^T$. $C_{j,t}$ is the ground-truth class for the j -th sounding object in the video at t frame. We search for an assignment $\sigma \in \mathcal{P}_N$ with the highest similarity where \mathcal{P}_N is a set of permutations of N elements. The similarity can be computed as

$$\mathcal{L}_{\text{match}}(y_i, \hat{y}_j) = \lambda_{\text{box}} \mathcal{L}_{\text{box}} + \lambda_{\text{cls}} \mathcal{L}_{\text{cls}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}}, \quad (10)$$

where λ_{box} , λ_{cls} , and λ_{mask} are weights to balance losses. We leverage a combination of Dice [35] and BCE loss as $\mathcal{L}_{\text{mask}}$, focal loss [40] as \mathcal{L}_{cls} , and GIoU [51] and L1 loss as \mathcal{L}_{box} . The best assignment $\hat{\sigma}$ is solved by the Hungarian algorithm [32]. Given the best assignment $\hat{\sigma}$, the segmentation loss between ground-truth and predictions is defined as $\mathcal{L}_{\text{segm}} = \mathcal{L}_{\text{match}}(y_i, \hat{y}_{\hat{\sigma}(j)})$.

4. Experiments

Dataset. We conduct experiments on AVS-Object [68] for AVS task and AVS-Semantic [69] for AVSS task.

Method	Backbone	AVS-Object-Single			AVS-Object-Multi			AVS-Semantic
		$\mathcal{J} \& \mathcal{F} \uparrow$	$\mathcal{J} \uparrow$	$\mathcal{F} \uparrow$	$\mathcal{J} \& \mathcal{F} \uparrow$	$\mathcal{J} \uparrow$	$\mathcal{F} \uparrow$	mIoU \uparrow
ResNet Backbone								
LVS [9]	ResNet-18	44.5	37.9	51.9	31.3	29.5	33.0	-
MSSL [50]	ResNet-18	55.6	44.9	66.3	31.4	26.1	36.3	-
3DC [42]	3DC	66.5	57.1	75.9	43.6	36.9	50.3	17.3
AOT [65]	ResNet-50	-	-	-	-	-	-	25.4
AVS [69]	ResNet-50	78.8	72.8	84.8	53.6	47.9	57.8	20.2
Bi-Gen [24]	ResNet-50	79.8	74.1	85.4	50.9	50.0	56.8	-
AVSegFormer [21]	ResNet-50	81.2	76.5	85.9	56.2	49.5	62.8	24.9
CATR [†] [34]	ResNet-50	81.0	74.9	87.1	59.4	53.1	65.6	-
Ours	ResNet-50	81.8	77.6	86.0	61.6	59.6	63.5	46.6
Transformer Backbone								
iGAN [43]	Swin-Tiny	69.7	61.6	77.8	48.7	42.9	54.4	-
SST [20]	SSL	73.2	66.3	80.1	49.9	42.6	57.2	-
LGVT [67]	Swin-Tiny	81.1	74.9	87.3	50.0	40.7	59.3	-
AVS [69]	PVT-v2-Base	83.3	78.7	87.9	59.3	54.0	64.5	29.8
Ours	Swin-Tiny	83.9	79.5	88.2	64.0	61.9	66.1	53.4

Table 1. **Quantitative comparison to AVS and AVSS methods.** \uparrow indicates the larger the better. \dagger denotes using additional training data.

- **AVS-Object:** AVS-Object dataset contains 5,356 short videos with corresponding audios in which 4,932 audios contain single-source and 424 audios contain multiple sources. Class-agnostic masks are given as annotations for AVS task. Typically, it is evaluated separately for single- and multi-source audios as **AVS-Object-Single** and **AVS-Object-Multi**.
- **AVS-Semantic:** AVS-Semantic is an extended dataset from AVS-Object which contains 12,356 videos with 70 classes. Semantic segmentation is annotated for AVSS task. Both single- and multi-source audio cases exist in the AVS-Semantic.

Metrics. For AVS task, the convention is to compute region similarity \mathcal{J} and contour accuracy \mathcal{F} as defined in [49]. Note that we follow the video segmentation convention to use the region similarity \mathcal{J} , which is equivalent to mIoU in the binary AVS setting. For AVSS, we follow the semantic segmentation convention to evaluate the model using mIoU which is defined as the intersection over union averaged among all classes.

Implementation Detail. We implement our method in PyTorch [47]. We train our model for 13 epochs and 16 epochs with a learning rate multiplier of 0.1 at the 11th and 14th epochs for AVS-Object and AVS-Semantic, respectively. We set the initial learning rate to be 1e-4, and a multiplier of 0.5 was applied to the backbone. We adopt batchsize = 4 and an AdamW [41] optimizer with weight decay 5×10^{-4} . The codebook size K is set to 128. The token number N for decomposing audio is set to 5. If no other specification, all images are resized to have the longest side 224 during evaluation. More details are available in the Appendix.

4.1. Main Results

Quantitative comparison on AVS-Object. We compare our method with existing AVS methods with CNN and transformer backbones. Without loss of generality, we adopt ResNet-50 and the small Swin-Tiny backbone, and provide results of larger backbones in the Appendix. Note that PVT-v2-Base is much larger than Swin-Tiny. Our method outperforms the previous state-of-the-art (SOTA) method AVSegFormer [21] and CATR [34] by 5.4 and 2.2 of $\mathcal{J} \& \mathcal{F}$ score on AVS-Object-Multi datasets respectively (with ResNet-50 backbone). The promising \mathcal{F} score achieved by CATR can be attributed to its additional incorporation of atrous spatial pyramid pooling [11], a feature that can be seamlessly integrated into our method as well. We notice that the improvement on the multi-source setting is much larger than the single-source setting. This is because single-source audios contain simple and disentangled semantics and can be easily aligned with visual features while, for multi-source audios, the complex semantic space makes the alignment to visual contents much more difficult.

Quantitative comparison on AVS-Semantic. Compared to the AVS-Object task, our method demonstrates greater improvement in the AVS-Semantic task. As shown in the Tab. 1, our method eclipses the previous SOTA AVSS method AOT [65] by a remarkable 21.2 mIoU with ResNet-50 backbone. The improvement in the AVSS task can be attributed to several factors. First, the task itself involves the semantic prediction of sound sources. However, due to mixed audio signals, aligning visual content accurately becomes challenging, leading to difficulties in classification. Secondly, the number of sound sources and categories of

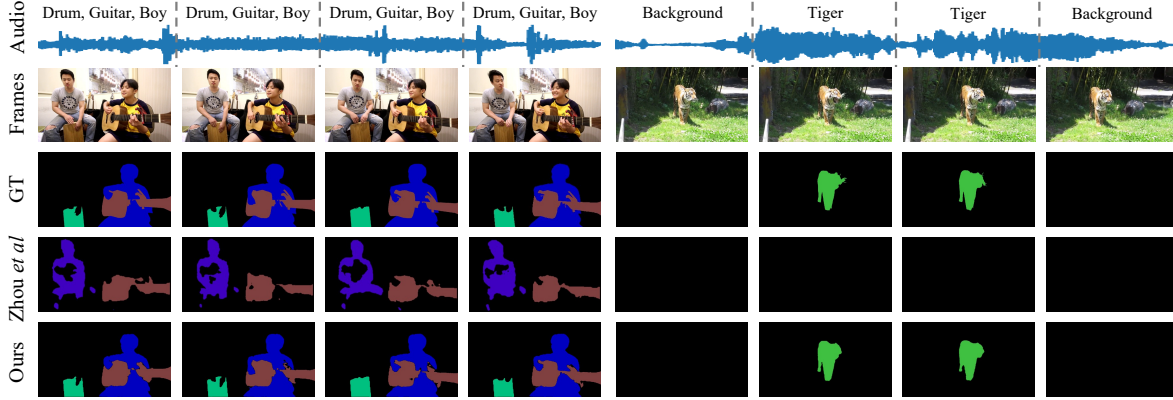


Figure 5. Qualitative comparison to Zhou et al. [69] on AVS-Semantic. Each color represents a semantic category. Note that we show the class labels in the first row as references but not given in the input, and audio signals and frames are the input.

ID	Method	AVS-Object-Multi			AVS-Semantic
		$\mathcal{J}\&\mathcal{F} \uparrow$	$\mathcal{J} \uparrow$	$\mathcal{F} \uparrow$	mIoU \uparrow
1	Baseline (w/o decomposition)	52.9 -8.7	50.1	55.7	33.5 -13.1
2	+ TD [10]	56.2 -5.4	54.5	57.9	38.9 -7.7
3	+ Non-Q-based-SD	57.6 -4.0	56.1	59.1	39.4 -6.8
4	+ Q-based-SD	59.7 -1.9	57.6	61.8	42.5 -4.5
5	+ Q-based-SD + AVSR	60.1 -1.5	58.2	61.9	44.5 -2.1
6	+ Q-based-SD + AVSR + LC (ours)	61.6	59.6	63.5	46.6

Table 2. **Component effectiveness study.** TD: decompose audio at the time domain in advance with an off-the-shelf model [10]; Non-Q-SD: semantic decomposition without quantization; Q-based-SD: quantization-based semantic decomposition; AVSR: audiovisual semantic recombination; LC: local calibration.

AVS-Semantic are larger than AVS-Object, which will result in a larger semantic space. When the mixed semantics are not decomposed, the network struggles to handle the numerous mixed semantics effectively. Thirdly, in the AVS-Semantic dataset, sound event changes occur more frequently. As a result, a more robust frame-level audio-visual correspondence is required. Our proposed global-to-local distilling mechanism addresses this challenge by enhancing the capture of local semantic information, enabling accurate object segmentation.

Qualitative comparison. As shown in Fig. 5, we qualitatively compare our method to the method proposed by Zhou et al. [68] on AVS-Semantic. Our method achieves better results on both segmenting quality and class prediction accuracy. Since the method [68] directly fuses mixed audio features with video features, we notice that it suffers from object incorrectness when multiple sound sources are present. Meanwhile, due to the lack of frame-level audio-visual calibration, [68] cannot effectively handle the audio semantic changes. More qualitative results are available in the Appendix and supplemental video.

4.2. Ablation Study

We carry out ablation studies on multi-source audio scenarios using the AVS-Object-Multi and AVS-Semantic

datasets. These studies aim to assess the effectiveness of various components and design choices in our proposed method for quantization-based semantic decomposition, as presented in Tab. 2. We use the ResNet-50 backbone for all ablation studies.

Effectiveness of audio decomposition. We establish a baseline model that includes only the two unimodal encoders, cross-attention for fusion, and the mask decoder. This model does not incorporate any decomposition but performs audiovisual interaction directly with cross-attention layers. The model’s performance significantly deteriorates in multi-source audio scenarios, yielding scores of 52.9 $\mathcal{J}\&\mathcal{F}$ for the AVS-Object-Multi dataset and 33.5 mIoU for the AVS-Semantic dataset (Row 1). This is an 8.7 $\mathcal{J}\&\mathcal{F}$ and 13.1 mIoU decrease compared to our semantic decomposition method (Row 6).

Effectiveness of decomposition on the semantic domain. As a straightforward method for audio decomposition is to separate audio signals in advance with pre-trained models, we provide an alternative method that first applies a widely used sound source separation model [10] and then performs audiovisual segmentation for each decomposed audio using our baseline model. This method has gained 3.3 $\mathcal{J}\&\mathcal{F}$ and 5.5 mIoU from the baseline, which also demonstrates that audio decomposition is helpful. However, it lags be-




Figure 6. Visualization comparison between entangled and our disentangled audio semantic space. “M” and “S” notations denote multi-source and single-source inputs.

hind our method that decomposes audio in the semantic domain, with a decrease of 5.4 $\mathcal{J}\&\mathcal{F}$ and 7.7 mIoU. We attribute this improvement to two factors: 1) the imperfection of sound source separation model, and 2) the conflicts that arise when combining the masks for each source in the time domain without considering visual content during separation. In contrast, our semantic-domain approach does not suffer from these issues and can effectively leverage the information contained in both audio and visual modalities.

Effectiveness of quantization for semantic decomposition. To validate the role of quantization, we set up a model that excludes the codebook learning and preserves all other modules for comparison. Without quantization, the results exhibit 4.0 $\mathcal{J}\&\mathcal{F}$ and 6.8 mIoU decrease (Row 3). The information bottleneck principle enhances our semantic decomposition by reducing irrelevant noise, clustering semantic-relevant features for disentanglement, and improving the matching and comparison of features in the segmentation task. In addition, for further fusing with visual features, we perform audiovisual semantic recombination (AVSR) to render the decomposed audio tokens aware of visual features. The results of Row 5 show an increase of 0.4 in $\mathcal{J}\&\mathcal{F}$ and 2.4 in mIoU (as is shown in Row 4).

Effectiveness of local calibration. We perform the local calibration by enforcing local quantization to align with global ones using a shared codebook and not updating it after local LSD. With the local calibration scheme, the audio representation is learned more robustly with long-term information, gaining 1.5 in $\mathcal{J}\&\mathcal{F}$ and 2.1 in mIoU (as is shown in Row 5).

Ablation on codebook size, semantic token number, frame number, and input resolution are in the Appendix.

4.3. Analysis

We delve deeper into our research by visualizing the feature space to understand our semantic decomposition better and exploring its robustness against background disturbances.

Visualization of decomposed semantic space. As shown in Fig. 6, we visualize the semantic space with and without semantic decomposition on the AVS-Semantic dataset using t-SNE [59]. Three types of single-source audios (“man”,




Figure 7. Analysis with different degrees of background noises (invisible sound sources) on AVSS task. The noise degree is measured by the signal-to-noise ratio (SNR) of the original sound.

“guitar”, “piano”) and two types of multi-source audios (“man+guitar”, “man+piano”) are enrolled. Without decomposition, the multi-source features are highly entangled, presenting less evidence related to single-source semantics. However, after performing semantic decomposition, the “man+guitar” feature presents clear evidence related to its corresponding single-source (“man” and “guitar”) semantics. This is reflected in the proximity of the “man+guitar” feature to the centroids of its corresponding single-source features. The same applies to the “man+piano” feature.

Analysis of robustness against background noises. We address the susceptibility of audio signals to background disturbances by employing quantized audio representation. In Fig. 7, we evaluate performance under varying degrees of background noises (drawn from AudioSet [22]). As the audio signal-to-noise ratio (SNR) decreases, the model utilizing quantized and decomposed audio representation consistently outperforms the one using conventional continuous representation. This superior performance can be attributed to two factors: (1) effective quantization discards unnecessary noisy details, enhancing audio representation robustness, and (2) PQ-based decomposition disentangles target audio semantics from noisy elements, reducing the impact of background noises in audiovisual interaction. In summary, our approach demonstrates improved resilience to background noise.

5. Conclusion

This paper explores the quantized and decomposed audio representation to address the audiovisual segmentation. Our approach effectively disentangles multi-source audio semantics with a PQ-based decomposition, providing a more robust audio representation for audiovisual interaction in complex environments. The global-to-local mechanism further enhances frame-level audio semantics, mitigating the instability associated with short-term extractions. Through extensive experiments, our approach has demonstrated superior performance compared to previous methods, particularly excelling in multi-object scenarios.

References

- [1] Alessandro Achille and Stefano Soatto. Emergence of invariance and disentanglement in deep representations. *The Journal of Machine Learning Research*, 19(1):1947–1980, 2018. 2
- [2] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 208–224. Springer, 2020. 3
- [3] Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016. 2
- [4] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017. 3
- [5] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)*, pages 435–451, 2018. 3
- [6] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multi-modal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4985–4995, 2022. 3
- [7] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multi-modal transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4985–4995, 2022. 1, 2
- [8] 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. 3
- [9] 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. 3, 6
- [10] Ke Chen, Xingjian Du, Bilei Zhu, Zejun Ma, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. Zero-shot audio source separation through query-based learning from weakly-labeled data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4441–4449, 2022. 7
- [11] 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, 2017. 6
- [12] Weidong Chen, Dexiang Hong, Yuankai Qi, Zhenjun Han, Shuhui Wang, Laiyun Qing, Qingming Huang, and Guorong Li. Multi-attention network for compressed video referring object segmentation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4416–4425, 2022. 3
- [13] Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. Dynamic convolution: Attention over convolution kernels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11030–11039, 2020. 5
- [14] Bowen Cheng, Anwesa Choudhuri, Ishan Misra, Alexander Kirillov, Rohit Girdhar, and Alexander G Schwing. Mask2former for video instance segmentation. *arXiv preprint arXiv:2112.10764*, 2021. 5
- [15] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. *arXiv preprint arXiv:2103.07941*, 2021. 3
- [16] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Re-thinking space-time networks with improved memory coverage for efficient video object segmentation. *arXiv preprint arXiv:2106.05210*, 2021. 3
- [17] Ying Cheng, Ruize Wang, Zhihao Pan, Rui Feng, and Yuejie Zhang. Look, listen, and attend: Co-attention network for self-supervised audio-visual representation learning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3884–3892, 2020. 3
- [18] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023. 1
- [19] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, and Chen Change Loy. Mevis: A large-scale benchmark for video segmentation with motion expressions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2694–2703, 2023. 3
- [20] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. *arXiv preprint arXiv:2101.08833*, 2021. 6
- [21] Shengyi Gao, Zhe Chen, Guo Chen, Wenhai Wang, and Tong Lu. Avsegformer: Audio-visual segmentation with transformer. *arXiv preprint arXiv:2307.01146*, 2023. 2, 6
- [22] Jort F Gemmeke, Daniel PW 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. IEEE, 2017. 8
- [23] Mingfei Han, Yali Wang, Zhihui Li, Lina Yao, Xiaojun Chang, and Yu Qiao. Htm: Hybrid temporal-scale multimodal learning framework for referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13414–13423, 2023. 3
- [24] 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. 6
- [25] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *2017 IEEE in-*

- ternational conference on acoustics, speech and signal processing (icassp), pages 131–135. IEEE, 2017. 3
- [26] Di Hu, Feiping Nie, and Xuelong Li. Deep multimodal clustering for unsupervised audiovisual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9248–9257, 2019. 3
- [27] Li Hu, Peng Zhang, Bang Zhang, Pan Pan, Yinghui Xu, and Rong Jin. Learning position and target consistency for memory-based video object segmentation. *arXiv preprint arXiv:2104.04329*, 2021. 3
- [28] Ping Hu, Fabian Caba, Oliver Wang, Zhe Lin, Stan Sclaroff, and Federico Perazzi. Temporally distributed networks for fast video semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8818–8827, 2020. 3
- [29] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pages 2117–2126. IEEE, 2017. 3
- [30] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1): 117–128, 2010. 2, 3
- [31] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 1
- [32] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [33] Dezhuang Li, Ruoqi Li, Lijun Wang, Yifan Wang, Jinqing Qi, Lu Zhang, Ting Liu, Qingquan Xu, and Huchuan Lu. You only infer once: Cross-modal meta-transfer for referring video object segmentation. In *AAAI Conference on Artificial Intelligence*, 2022. 3
- [34] 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
- [35] Xiaoya Li, Xiaofei Sun, Yuxian Meng, Junjun Liang, Fei Wu, and Jiwei Li. Dice loss for data-imbalanced nlp tasks. *arXiv preprint arXiv:1911.02855*, 2019. 5
- [36] Xiang Li, Jinglu Wang, Xiaohao Xu, Xiao Li, Bhiksha Raj, and Yan Lu. Robust referring video object segmentation with cyclic structural consensus. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22236–22245, 2023. 3
- [37] Yule Li, Jianping Shi, and Dahua Lin. Low-latency video semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5997–6005, 2018. 3
- [38] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 3
- [39] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 5
- [40] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [41] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [42] Sabarinath Mahadevan, Ali Athar, Aljoša Ošep, Sebastian Hennen, Laura Leal-Taixé, and Bastian Leibe. Making a case for 3d convolutions for object segmentation in videos. *arXiv preprint arXiv:2008.11516*, 2020. 6
- [43] Yuxin Mao, Jing Zhang, Zhexiong Wan, Yuchao Dai, Aixuan Li, Yunqiu Lv, Xinyu Tian, Deng-Ping Fan, and Nick Barnes. Transformer transforms salient object detection and camouflaged object detection. *arXiv preprint arXiv:2104.10127*, 2021. 6
- [44] Bo Miao, Mohammed Bennamoun, Yongsheng Gao, and Ajmal Mian. Spectrum-guided multi-granularity referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 920–930, 2023. 3
- [45] Shentong Mo and Pedro Morgado. A closer look at weakly-supervised audio-visual source localization. *arXiv preprint arXiv:2209.09634*, 2022. 3
- [46] Shentong Mo and Pedro Morgado. Localizing visual sounds the easy way. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII*, pages 218–234. Springer, 2022. 3
- [47] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019. 6
- [48] Matthieu Paul, Christoph Mayer, Luc Van Gool, and Radu Timofte. Efficient video semantic segmentation with labels propagation and refinement. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2873–2882, 2020. 3
- [49] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 6
- [50] 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: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 292–308. Springer, 2020. 3, 6
- [51] Hamid Rezaatoughi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. 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*, pages 658–666, 2019. 5

- [52] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4358–4366, 2018. 3
- [53] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *European Conference on Computer Vision*, pages 208–223. Springer, 2020. 3
- [54] Hongje Seong, Junhyuk Hyun, and Euntai Kim. Kernelized memory network for video object segmentation. In *European Conference on Computer Vision*, pages 629–645. Springer, 2020. 3
- [55] Hongje Seong, Seoung Wug Oh, Joon-Young Lee, Seongwon Lee, Suhyeon Lee, and Euntai Kim. Hierarchical memory matching network for video object segmentation, 2021. 3
- [56] Guolei Sun, Yun Liu, Henghui Ding, Thomas Probst, and Luc Van Gool. Coarse-to-fine feature mining for video semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3126–3137, 2022. 3
- [57] Jiajin Tang, Ge Zheng, and Sibe Yang. Temporal collection and distribution for referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15466–15476, 2023. 3
- [58] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. 2
- [59] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (11), 2008. 8
- [60] Dongming Wu, Xingping Dong, Ling Shao, and Jianbing Shen. Multi-level representation learning with semantic alignment for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4996–5005, 2022. 3
- [61] Dongming Wu, Tiancai Wang, Yuang Zhang, Xiangyu Zhang, and Jianbing Shen. Onlinerefer: A simple online baseline for referring video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2761–2770, 2023. 1
- [62] Jiannan Wu, Yi Jiang, Peize Sun, Zehuan Yuan, and Ping Luo. Language as queries for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4974–4984, 2022. 2, 3
- [63] Jiannan Wu, Yi Jiang, Peize Sun, Zehuan Yuan, and Ping Luo. Language as queries for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4974–4984, 2022. 1
- [64] Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, and Aggelos K Katsaggelos. Efficient video object segmentation via network modulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6499–6507, 2018. 3
- [65] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. 3, 6
- [66] Chen Yuanhong, Liu Yuyuan, Wang Hu, Liu Fengbei, Wang Chong, and Carneiro Gustavo. A closer look at audio-visual semantic segmentation. *arXiv preprint arXiv:2304.02970*, 2023. 3
- [67] Jing Zhang, Jianwen Xie, Nick Barnes, and Ping Li. Learning generative vision transformer with energy-based latent space for saliency prediction. *Advances in Neural Information Processing Systems*, 34:15448–15463, 2021. 6
- [68] 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*, 2022. 1, 2, 3, 5, 7
- [69] 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, 2, 3, 5, 6, 7
- [70] Jiafan Zhuang, Zilei Wang, and Yuan Gao. Semi-supervised video semantic segmentation with inter-frame feature reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3263–3271, 2022. 3