

# Discovering Sounding Objects by Audio Queries for Audio Visual Segmentation

Shaofei Huang<sup>1,2</sup>, Han Li<sup>3</sup>, Yuqing Wang<sup>4</sup>, Hongji Zhu<sup>4</sup>,  
 Jiao Dai<sup>1,2\*</sup>, Jizhong Han<sup>1,2</sup>, Wenge Rong<sup>3</sup> and Si Liu<sup>5,6</sup>

<sup>1</sup>Institute of Information Engineering, Chinese Academy of Sciences

<sup>2</sup>School of Cyber Security, University of Chinese Academy of Sciences

<sup>3</sup>School of Computer Science and Engineering, Beihang University

<sup>4</sup>Alibaba Group

<sup>5</sup>Institute of Artificial Intelligence, Beihang University

<sup>6</sup>Hangzhou Innovation Institute, Beihang University

## Abstract

Audio visual segmentation (AVS) aims to segment the sounding objects for each frame of a given video. To distinguish the sounding objects from silent ones, both audio-visual semantic correspondence and temporal interaction are required. The previous method applies multi-frame cross-modal attention to conduct pixel-level interactions between audio features and visual features of multiple frames simultaneously, which is both redundant and implicit. In this paper, we propose an Audio-Queried Transformer architecture, AQFormer, where we define a set of object queries conditioned on audio information and associate each of them to particular sounding objects. Explicit object-level semantic correspondence between audio and visual modalities is established by gathering object information from visual features with predefined audio queries. Besides, an Audio-Bridged Temporal Interaction module is proposed to exchange sounding object-relevant information among multiple frames with the bridge of audio features. Extensive experiments are conducted on two AVS benchmarks to show that our method achieves state-of-the-art performances, especially 7.1%  $M_J$  and 7.6%  $M_F$  gains on the MS3 setting.

## 1 Introduction

It is intuitive that objects are characterized not only by what they look like but also by the sounds they make. By excavating the correspondence between visual and acoustic information of objects, it is helpful to facilitate the understanding of their characteristics. In this work, we focus on the problem of audio visual segmentation (AVS) [Zhou *et al.*, 2022], which aims to identify and segment the sounding objects in each frame of the given video. Unlike audio-visual correspondence [Arandjelovic and Zisserman, 2017] or sound source localization [Arandjelovic and Zisserman, 2018] that only produce the interval of audible frames or the heatmap

\*Corresponding author

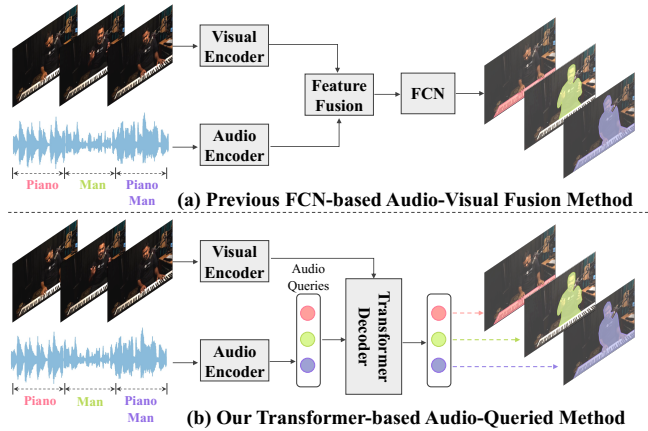


Figure 1: Comparison between different methods. (a) The previous method establishes pixel-level audio-visual correspondence by fusing audio and visual features with cross-modal attention. (b) Our method establishes object-level audio-visual correspondence using audio queries to gather object information from visual features.

centered on the sounding objects, AVS requires more fine-grained correspondence between audio and visual modalities. In addition, the situation of sounding in each video clip is dynamically changing over time. For example, different objects may make sounds at different time periods, and multiple objects may make sounds at the same time period. The uncertainty of target objects also makes AVS even more challenging than other video object segmentation tasks [Perazzi *et al.*, 2016; Seo *et al.*, 2020] where the target object remains the same for the whole video.

There exist two challenges in identifying the sounding objects under complex audio-visual scenes. On the one hand, objects of different categories typically differ from each other in the sounds they emit, indicating that we can distinguish them by their acoustic characteristics. As illustrated in Figure 1, there is an obvious difference between the waveform diagrams of the human voice and the guitar sound. Thus, it is critical to establish proper semantic correspondence between audio and visual features. On the other hand, due to the similarity of acoustic characteristics between objects of the same

category, it is difficult to distinguish them with their audio features, which may lead to confusion between them. For example, there are two men in the video of Figure 1 who sing in the first frame and second frame respectively, making it hard to tell them apart without additional clues. Therefore, temporal contexts such as action and mouth movement can be exploited to locate the sounding objects in such cases.

As shown in Figure 1(a), the previous method [Zhou *et al.*, 2022] first applies a multi-frame cross-modal attention module to fuse the visual features and the audio features extracted by a visual encoder and an audio encoder respectively, and then makes predictions on the fused multi-modal features with an FCN head. It has some limitations in dealing with the above two challenges. First, the cross-modal attention interacts between visual and audio features only at the pixel level, which is too implicit and may include unwanted background pixels in the images. Such pixel-level audio-visual semantic correspondence leads to an insufficient understanding of the sounding characteristics of different objects. Second, the temporal information modeling is realized in a mixed manner by the attention module, where each pixel of the video frames densely interacts with all pixels through the attention mechanism. This way of dense temporal interaction tends to be inefficient and redundant in capturing the temporal movements of sounding objects.

In this work, we propose a new Audio-Queried Transformer architecture, **AQFormer**, to tackle the above two challenges. *For the first one*, we define a set of audio-conditioned object queries to build explicit object-level correspondence between audio and visual modalities. As illustrated in Figure 1(b), each audio query is generated by the audio feature of one frame and is associated with one or more objects that make sounds in the corresponding frame but may also exist in other frames. By gathering object information from the visual features, the audio queries gradually encode the global information of the associated sounding objects, so that the generated object embeddings can be used to produce masks for the corresponding frames. An auxiliary loss is also designed to constrain the object-level feature similarity between the audio queries and visual features, which helps to learn more discriminative object representations. *For the second one*, in order to model the temporal pattern of the sounding objects, we enable each audio query to gather object information from the visual features of all the video frames. In this way, the temporal information is integrated into the object embeddings to facilitate the distinguishment between visually confusing objects. In addition, to further endow the visual features with sounding object-aware temporal context, we propose an Audio-Bridged Temporal Interaction (ABTI) module which leverages the audio feature as a medium to bridge the interaction between different frames. Concretely, the audio features are first utilized to gather the audio-relevant visual features of each frame. The compact features of different frames interact with each other to exchange information in a denoised manner and are then remapped back to the corresponding visual feature to enhance it with temporal context. The ABTI module is densely inserted after each stage of the visual encoder to generate temporal-aware visual features.

Our contributions are summarized as follows:

(1) We propose a novel Audio-Queried Transformer Architecture, AQFormer, for audio visual segmentation. A set of object queries conditioned on audio information are defined and utilized to gather visual information of their associated sounding objects for mask prediction.

(2) We design an Audio-Bridged Temporal Interaction (ABTI) module to exchange sounding object-relevant information among multiple frames bridged by audio features.

(3) Extensive experiments show that our AQFormer significantly outperforms the previous FCN-based method and achieves state-of-the-art performances on both single-source and multi-source AVS benchmarks.

## 2 Related Works

### 2.1 Video Segmentation

**Video Semantic Segmentation.** Generally, the popular convolutional-based segmentation methods [Long *et al.*, 2015] make pixel-wise classification to classify each pixel into a predefined category by a fully convolutional layer. Later works make efforts to capture multi-scale context information [Chen *et al.*, 2017] or aggregate features with long-range dependencies [Fu *et al.*, 2019; Yuan *et al.*, 2021]. With the development of vision transformer [Dosovitskiy *et al.*, 2020; Strudel *et al.*, 2021] and object query methods [Carion *et al.*, 2020], MaskFormer [Cheng *et al.*, 2021] first formulates the segmentation problem as mask classification and defines mask queries to encode region features for mask prediction. Mask2Former [Cheng *et al.*, 2022] further limits the query scopes to local regions by proposing masked attention. When extended to video semantic segmentation, most methods focus on the problems of temporal information extraction and inter-frame relationship modeling [Nilsson and Sminchisescu, 2018; Zhu *et al.*, 2017]. STFCN [Fayyaz *et al.*, 2016] introduces LSTM [Hochreiter and Schmidhuber, 1997] to integrate features from multiple frames sequentially. Similarly, STGRU [Nilsson and Sminchisescu, 2018] inserts GRU [Cho *et al.*, 2014] in the pipeline to control the feature fusion at different timesteps.

**Video Object Segmentation.** Video object segmentation (VOS) aims to track and segment foreground objects of a video in a class-agnostic manner. Generally, additional information is provided as a hint to indicate the target object to segment in VOS tasks [Yang and Yang, 2021; Yang and Yang, 2022]. For example, the object mask of the first frame, i.e., the reference frame, is provided in semi-supervised video object segmentation. To propagate object masks from the reference frame, STM [Oh *et al.*, 2019] introduces a space-time memory network that collects information from both the reference frame and the previous frame. CFBI [Yang *et al.*, 2020] extracts features from both foreground objects and background regions and matches them with the first frame to refine object masks. Another conditional VOS task, Referring Video Object Segmentation (RVOS) [Seo *et al.*, 2020] provides the linguistic expression that describes the object’s characteristics to indicate the target object. LBDT [Ding *et al.*, 2022] proposes a Language-Bridged Duplex Transfer (LBDT) module which utilizes language as an intermediary bridge to exchange information be-

tween spatial and temporal branches. ReferFormer [Wu *et al.*, 2022b] and MTTR [Botach *et al.*, 2022] both adopt multimodal transformer architectures that use text embeddings to restrict the object queries and predict object sequence by linking objects from different frames.

### 2.2 Audio-Visual Learning

Audio and visual signals usually co-occur in the same video, which provides rich semantic information for effective multimodal learning and scene perception. Some works [Arandjelovic and Zisserman, 2017; Wu *et al.*, 2022a; Alwassel *et al.*, 2020] leverage the natural correspondence between audio and visual information as supervision to learn discriminative multi-modal feature representations. These methods are generally utilized for pretraining and only segment-level audio-visual correspondence is constructed. Different from the above methods, Sound Source Localization (SSL) [Arandjelovic and Zisserman, 2018; Zhao *et al.*, 2019; Rouditchenko *et al.*, 2019] leverages the relevance between audio and vision regions to distinguish different sound sources and locate audio-related sounding areas. For example, DMC [Hu *et al.*, 2019] clusters audio and visual representations within each modality and associates the centroids of both modalities with contrastive learning. A two-stage framework is proposed in [Qian *et al.*, 2020] to align features of audio-visual modalities from coarse to fine. SSL only learns patch-level audio-visual correspondence by producing a heat map of the sounding objects, which is insufficient for applications that require more fine-grained visual results.

In this work, we focus on the problem of audio-visual segmentation (AVS) [Zhou *et al.*, 2022], which further requires to delineate the shape of the sounding objects at each frame. A baseline method TPAVI is also proposed in [Zhou *et al.*, 2022] to fuse audio features with video features by temporal pixel-wise audio-visual attention mechanism. Compared with SSL, AVS achieves pixel-level alignment between visual and audio modalities, enabling more comprehensive scene understanding.

## 3 Method

Given a video clip with  $T$  frames which is represented by its corresponding visual and audio components, our goal is to obtain the binary mask of sounding objects for each frame. The overall architecture of our proposed AQFormer is illustrated in Figure 2. A visual encoder is first adopted to extract multi-scale visual features for each frame. An audio encoder is used to process the input audio signal and obtain the corresponding audio feature. Then, we feed both the multi-scale visual features and the audio feature into the pixel decoder to obtain the audio-relevant and temporally-enhanced visual features. Afterward, the first three scales of enhanced visual features are fed into the audio-queried transformer decoder, in which the audio feature is used to generate  $T$  conditional object queries to collect object features from the enhanced features, and object embeddings are produced for each frame. The final predictions are obtained via matrix multiplication between the object embeddings and the largest visual feature output by the pixel encoder.

### 3.1 Visual and Audio Encoders

Given a video clip with  $T$  frames, we adopt ResNet-50 [He *et al.*, 2016] or PVT-v2 [Wang *et al.*, 2022] as the visual encoder to extract the visual features for each frame separately following [Zhou *et al.*, 2022]. To capture multi-scale visual context, we extract visual features of 4 different scales, including 1/32, 1/16, 1/8, and 1/4 of the original image size. The visual feature of the  $i$ -th scale is denoted as  $\mathbf{V}^i \in \mathbb{R}^{T \times H^i \times W^i \times C_v^i}$ , where  $H^i, W^i$  denote the height and width of the  $i$ -th feature and  $C_v^i$  denotes its channel number. As for the audio input, we first process it with the Short-time Fourier Transform to generate its spectrogram. A convolutional network, such as VGGish [Hershey *et al.*, 2017] is then applied to the spectrogram to extract the audio feature  $\mathbf{A} \in \mathbb{R}^{T \times C_a}$ , where  $C_a$  denotes the channel number of  $\mathbf{A}$ . In the following sections, we leverage subscript  $t$  to index audio or visual feature of the  $t$ -th frame, e.g.,  $\mathbf{A}_t \in \mathbb{R}^{C_a}$  and  $\mathbf{V}_t^i \in \mathbb{R}^{H^i \times W^i \times C_v^i}$ .

### 3.2 Pixel Decoder

In pixel decoder, the multi-scale visual features are enhanced to be aware of both the audio information and temporal context. First, the audio feature is integrated to enhance the corresponding frame of visual feature by cross-attention. Since operations performed on each scale of visual features are the same, we take the  $i$ -th scale as an example and omit the superscript  $i$  for simplicity. The cross-attention is calculated as follows:

$$\bar{\mathbf{V}}_t = \text{Softmax}\left(\frac{f_q(\mathbf{V}_t) \otimes f_k(\mathbf{A}_t)^T}{\sqrt{C_m}}\right) \otimes f_v(\mathbf{A}_t), \quad (1)$$

$$\mathbf{M}_t = \bar{\mathbf{V}}_t + f_w(\mathbf{V}_t), \quad (2)$$

where  $\mathbf{A}_t$  and  $\mathbf{V}_t$  are first reshaped to  $\mathbb{R}^{1 \times C_a}$  and  $\mathbb{R}^{HW \times C_v}$  respectively,  $\otimes$  denotes matrix multiplication, and  $f_q(\cdot), f_k(\cdot), f_v(\cdot)$  and  $f_w(\cdot)$  denote linear transform operations to change the channel numbers of input features to  $C_m, C_m, C_f$  and  $C_f$  respectively. Following [Cheng *et al.*, 2022], the enhanced visual feature  $\mathbf{M}^i$  is then reshaped to  $\mathbb{R}^{T \times H^i \times W^i \times C_f}$  and fed into the multi-scale deformable attention [Zhu *et al.*, 2021] module to obtain the aggregated feature  $\mathbf{M}^i$ . In order to reduce the computational cost, the above operations are performed on all scales of visual features except for the largest one. Since the above operations are performed on each frame independently without introducing temporal information, we further propose the Audio Bridged Temporal Interaction (ABTI) module to bridge the cross-frame interaction by audio feature. As shown in Figure 2, ABTI is applied to all scales of visual features, including the largest one  $\mathbf{V}^4$  and other aggregated features  $\mathbf{M}^i, i \in [1, 3]$ . The output features of pixel encoder are represented by  $\hat{\mathbf{M}}^i \in \mathbb{R}^{T \times H^i \times W^i \times C_f}, i \in [1, 4]$ . The details of ABTI are elaborated in Section 3.4.

### 3.3 Audio-Queried Transformer Decoder

In the audio-queried transformer decoder, we first define an object query for each frame to refer to the sounding objects in the current frame. For example, in Figure 2, the query



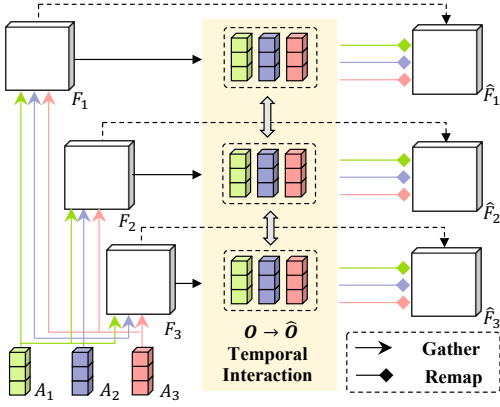


Figure 3: Illustration of ABTI module. Audio features are first leveraged to gather audio-relevant object features from different frames. After exchanging information between these object features, the updated object features are then remapped and added to the original features for enhancement.

features and we take one of them as an example for detailed elaboration, which is denoted as  $\mathbf{F} \in \mathbb{R}^{T \times H \times W \times C_v}$  in this section only.  $H$  and  $W$  represent the height and width of  $\mathbf{F}$ . As illustrated in Figure 3, there are three steps in the ABTI module, including object feature gathering, temporal interaction, and object feature remapping. First, for the sounding objects associated with the  $p$ -th audio feature, we extract its compact feature at the  $q$ -th frame from  $\mathbf{F}_q$  by feature similarity as follows:

$$\mathbf{S}_{pq} = \text{Softmax}(f_a(\mathbf{A}_p) \otimes f_m(\mathbf{F}_q)^T), \quad (9)$$

$$\mathbf{O}_{pq} = \mathbf{S}_{pq} \otimes f_n(\mathbf{F}_q), \quad (10)$$

where  $\mathbf{A}_p$  and  $\mathbf{F}_q$  are first reshaped to  $\mathbb{R}^{1 \times C_v}$  and  $\mathbb{R}^{H \times W \times C_a}$  respectively,  $f_a(\cdot)$ ,  $f_m(\cdot)$ , and  $f_n(\cdot)$  change the channel numbers of input features to  $C_b$  by linear transformation. Each element of  $\mathbf{S}_{pq} \in \mathbb{R}^{1 \times H \times W}$  represents the contribution of the corresponding pixel to the sounding objects, and the compact feature  $\mathbf{O}_{pq} \in \mathbb{R}^{C_b}$  is generated through necessary reshaping operation. We arrange the compact features of all the sounding objects at all the frames as a sequence  $\mathbf{O} = \{\mathbf{O}_{pq}\}, p \in [1, T], q \in [1, T]$ . Second, self-attention is applied to  $\mathbf{O}$  for temporal interaction and the updated object features are represented as  $\hat{\mathbf{O}} = \{\hat{\mathbf{O}}_{pq}\}, p \in [1, T], q \in [1, T]$ . Last, to enhance the original feature  $\mathbf{F}_q$  with temporal information from other frames, the updated object features are remapped back to match the shape of  $\mathbf{F}_q$  as follows:

$$\tilde{\mathbf{F}}_{pq} = \mathbf{S}_{pq}^T \otimes f_o(\hat{\mathbf{O}}_{pq}), \quad (11)$$

where  $\hat{\mathbf{O}}_{pq}$  is first reshaped to  $\mathbb{R}^{1 \times C_b}$  and then transformed by  $f_o(\cdot)$  to the channel number of  $C_f$ . After change the shape of  $\tilde{\mathbf{F}}_{pq}$  to  $\mathbb{R}^{H \times W \times C_v}$ , the enhanced feature  $\hat{\mathbf{F}}_q$  is generated by as follows:

$$\hat{\mathbf{F}}_q = \mathbf{F}_q + \sum_{p=1}^T \tilde{\mathbf{F}}_{pq}. \quad (12)$$

The above operations are equally performed on all the visual features to enable cross-frame interaction.

### 3.5 Loss Functions

The overall loss function is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{mask} + \lambda_{sim} \mathcal{L}_{sim}, \quad (13)$$

where  $\lambda_{mask}$  consists of Binary Cross-Entropy Loss and Dice Loss [Milletari *et al.*, 2016] following the implementation of [Cheng *et al.*, 2022]:

$$\lambda_{mask} = \lambda_{bce} \mathcal{L}_{bce} + \lambda_{dice} \mathcal{L}_{dice}, \quad (14)$$

the  $\mathcal{L}_{bce}$  and  $\mathcal{L}_{dice}$  are equally applied between the output of each decoder layer and the ground-truth mask as supervision.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

We conduct experiments on two benchmark settings of the AVS [Zhou *et al.*, 2022]: 1) The semi-supervised Single Sound Source Segmentation (S4) where the sounding object remains the same in the given video clip, i.e., only the mask annotation of the first frame is provided for training; 2) The fully supervised Multiple Sound Source Segmentation (MS3) where the sounding object dynamically changes over time, and mask annotations of all the  $T$  frames are provided. F-score  $M_{\mathcal{F}}$  and Jaccard index  $M_{\mathcal{J}}$  are adopted as evaluation metrics.  $M_{\mathcal{F}}$  calculates the harmonic mean of pixel-level precision and recall, while  $M_{\mathcal{J}}$  calculates the intersection-over-union between the predicted mask and the ground truth mask.

### 4.2 Implementation Details

We use both ResNet-50 [He *et al.*, 2016] pretrained on MSCOCO [Lin *et al.*, 2014] dataset and the PVT-v2 b5 [Wang *et al.*, 2022] pretrained on ImageNet [Russakovsky *et al.*, 2015] dataset as the visual encoders. The total number of video frames  $T$  is set to 5 for each video clip. The number  $N$  of transformer decoder stage is set 3.  $\lambda_{bce}$ ,  $\lambda_{dice}$  and  $\lambda_{sim}$  are all set to 1. For the S4 setting, we use the polynomial learning rate schedule and the AdamW optimizer with an initial learning rate of  $1.25e^{-4}$  and weight decay of  $5e^{-2}$ . Batch size is set to 8/4 and the total number of training iterations is set to 20k/40k for experiments on ResNet-50/PVT-v2 b5. The MS3 setting adopts the same hyperparameters except: (1) the initial learning rate is  $5e^{-4}$ , (2) the total training iterations are reduced to 2k/4k. More implementation details are provided in the supplementary materials.

### 4.3 Comparison with State-of-the-art Methods

Since AVS is a pretty new task and there is only one AVS method TPAVI [Zhou *et al.*, 2022] to compare at present, we follow [Zhou *et al.*, 2022] to adapt methods of related domains to AVS and present the results in Table 1. Taking ResNet-50 as the visual encoder, our AQFormer outperforms TPAVI by margins of 1.6%  $M_{\mathcal{F}}$  and 4.2%  $M_{\mathcal{J}}$  for the S4 setting. For the more challenging MS3 setting, AQFormer achieves 9.1%  $M_{\mathcal{F}}$  and 7.8%  $M_{\mathcal{J}}$  absolute gains over TPAVI, which is a notable improvement. The results indicate that the explicit object-level audio-visual correspondence construction leads to more robust performance when multiple objects make sounds at the same time or the sounding objects transfer. When using a stronger visual encoder like PVT-v2 b5,

Method	Task	Backbone	S4		MS3		MS3*	
			$M_{\mathcal{F}}$	$M_{\mathcal{J}}$	$M_{\mathcal{F}}$	$M_{\mathcal{J}}$	$M_{\mathcal{F}}$	$M_{\mathcal{J}}$
MSSL [Qian <i>et al.</i> , 2020]	SSL	ResNet-18	0.663	0.449	0.363	0.261	-	-
SST [Duke <i>et al.</i> , 2021]	VOS	ResNet-101	0.801	0.663	0.572	0.426	-	-
TPAVI [Zhou <i>et al.</i> , 2022]	AVS	ResNet-50	0.848	0.728	0.578	0.479	-	0.543
		PVT-v2	0.879	0.787	0.645	0.540	-	0.573
AQFormer(ours)	AVS	ResNet-50	<b>0.864</b>	<b>0.770</b>	<b>0.669</b>	<b>0.557</b>	<b>0.699</b>	<b>0.593</b>
		PVT-v2	<b>0.894</b>	<b>0.816</b>	<b>0.721</b>	<b>0.611</b>	<b>0.727</b>	<b>0.622</b>

Table 1: **Comparison with state-of-the-art methods.** Results of both S4 and MS3 settings are reported. MS3\* indicates that the model was first trained on the S4 training data.

Method	$M_{\mathcal{F}}$	$M_{\mathcal{J}}$
Vanilla Object Query	0.534	0.450
Audio Query†	0.643	0.526
<b>Audio Query</b>	<b>0.647</b>	<b>0.535</b>

(a) **Effect of audio query.** Audio Query† means each audio query only gathers object information from a single frame.

Inserting Positions	-	1,2,3	4	1,2,3,4
$M_{\mathcal{F}}$	0.650	0.652	0.652	<b>0.669</b>
$M_{\mathcal{J}}$	0.546	0.549	0.554	<b>0.557</b>

(c) **Inserting positions of ABTI.**

ABTI	AVSim Loss	$M_{\mathcal{F}}$	$M_{\mathcal{J}}$
		0.647	0.535
✓		0.656	0.543
	✓	0.650	0.546
✓	✓	<b>0.669</b>	<b>0.557</b>

(b) **Component analysis.**

Number of Stages	1	2	3
$M_{\mathcal{F}}$	0.599	0.632	<b>0.669</b>
$M_{\mathcal{J}}$	0.500	0.520	<b>0.557</b>

(d) **Number of transformer decoder stages.**

Table 2: **Ablation Studies.** Models are trained on training set of MS3 and evaluated on its testing set with ResNet-50 as the visual encoder.

further performance gains can be observed for both settings. We also conduct experiments on MS3 by using the model pre-trained on S4 for initialization, which also has improvements compared with training from scratch.

#### 4.4 Ablation Studies

We conduct all the ablation studies on MS3 benchmark with ResNet-50 as the visual encoder to evaluate the effectiveness of our method. More results are presented in the supplementary materials.

**Effect of audio query.** To verify the superiority of conditioning object queries on audio features over vanilla object queries, we replace audio queries with randomly initialized object queries in the transformer decoder and keep other settings the same for comparison, which is represented as ‘Vanilla Object Query’ in Table 2a. The large performance gap between vanilla object query and audio query indicates that gathering audio-relevant object information by audio queries can better identify the sounding objects for AVS. Besides, we also revise the cross attention in decoder layers so that each audio query only gathers object information from its corresponding frame rather than the whole video, which is denoted by Audio Query†. Performance drops can be observed on both  $M_{\mathcal{F}}$  and  $M_{\mathcal{J}}$  due to the absence of temporal information, indicating the effectiveness of temporal context gathering.

**Component Analysis.** We conduct ablation studies to explore the impact of ABTI module and the auxiliary loss,

which is denoted as ‘AVSim Loss’ in Table 2b. The first row denotes the basic implementation of our AQFormer, which has already outperformed the previous state-of-the-art by large margins ( $M_{\mathcal{F}} + 6.9%$ ,  $M_{\mathcal{J}} + 5.6%$ ), which shows the superiority of our new architecture. As shown in the 2nd and 3rd rows, both ABTI module and AVSim loss bring improvement over the baseline, i.e.,  $M_{\mathcal{F}} + 0.9%$  and  $M_{\mathcal{J}} + 0.8%$  for ABTI, and  $M_{\mathcal{F}} + 0.3%$  and  $M_{\mathcal{J}} + 1.1%$  for AVSim loss. The performance improvement indicates both temporal interaction and audio-visual feature alignment are necessary for AVS task. When combining the ABTI module and AVSim loss together in the 4th row, we can observe an obvious performance boost in all metrics compared with our baseline.

**Inserting Positions of ABTI.** We evaluate different inserting positions of the ABTI module in the pixel decoder and report the experimental results in Table 2c. It is found that larger gains on  $M_{\mathcal{J}}$  are obtained by applying ABTI to the 4th visual feature that serves as the mask feature for prediction. Applying ABTI to the first 3 scales of visual features achieves relatively slight improvement, which may be because that temporal information has already been modeled in the cross-attention steps of the transformer decoder. The best results are achieved when ABTI is added to all four stages, showing that temporal interaction still helps to get better performances.

**Number of Transformer Decoder Layers.** We also adjust the number of transformer stages  $N$  and the results are presented in Table 2d. By increasing the number of transformer




Figure 4: Qualitative comparison with previous method. For each video example, we sample 3 different frames for visualization and arrange them from left to right.

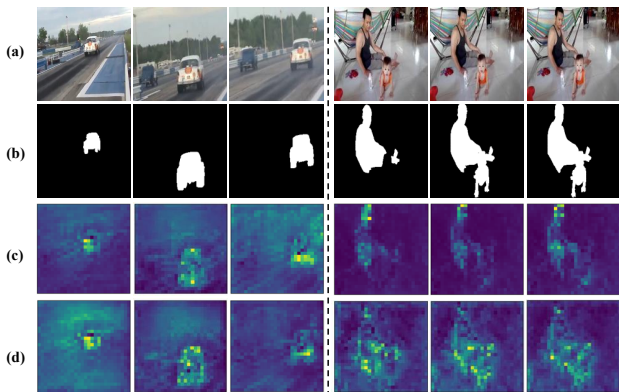


Figure 5: Visualization of attention maps. (a) Raw images of different frames. (b) Ground truth masks. (c) Attention maps of different frames attended by the first audio query. (d) Attention maps of different frames attended by the last audio query. Yellow denotes a larger attention weight. *Better viewed in color.*

stages from 1 to 3, we can observe consistent performance improvement. We get the best results when  $N$  is set to 3,  $0.669 M_{\mathcal{F}}$  and  $0.557 M_{\mathcal{J}}$ , which indicates that deeper transformer decoder results in better performances.

### 4.5 Qualitative Analysis

We compare the qualitative results between TPAVI [Zhou *et al.*, 2022] and our AQFormer in Figure 4. PVT-v2 is used as the visual encoder of both methods. Our AQFormer makes better identification and produces finer segmentation results as well. As shown in the 2nd and 3rd columns of the 1st example, when there are more helicopters appearing in the image but not making any sound, TPAVI cannot figure out the correct sounding objects and includes all of them. However, our AQFormer can accurately track the movement of the

sounding helicopter without generating false positive predictions. In the 3rd example, when the sounding objects increase from man only to man and baby, our AQFormer successfully captures the change while TPAVI does not include the baby as the sounding object in the last two frames, which further proves that audio query establishes more robust semantic correspondence between audio and visual signals.

We also visualize the attention maps between the audio queries and the visual features of the last layer of the transformer decoder. As illustrated in Figure 5, for each audio query, its attention maps at different frames always focus on the regions of the same object(s), even though its associated objects do not make sounds in some frames. Take the 4th row of the 2nd video as an example, the corresponding audio query is associated with the man and the baby, and only the man makes sound in the first frame. However, the activated pixels mainly focus on both the man and the baby in the first frame. These results show that our AQFormer indeed establishes an object-level correlation between audio and video modalities.

## 5 Conclusion

In this paper, we focus on the task of audio visual segmentation (AVS) and propose a new multi-modal transformer framework named Audio-Queried Transformer (AQFormer). Audio features are utilized to generate conditional object queries to explicitly gather information of the sounding objects from visual features. An Audio-Bridged Temporal Interaction module is also proposed to interact among multiple frames with the bridging of audio information. Extensive experiments on two AVS benchmarks show that our method outperforms the previous method by large margins. In the future, we hope to explore audio separation techniques and construct more specific instance-level correspondence between audio and visual information.

## Acknowledgements

This research was supported by National Key R&D Program of China (2022ZD0115502), National Natural Science Foundation of China (No. 62122010), Zhejiang Provincial Natural Science Foundation of China under Grant No. LDT23F02022F02, and Key Research and Development Program of Zhejiang Province (No. 2022C01082).

## Contribution Statement

**Shaofei Huang**<sup>†</sup>: Designing the architecture of AQFormer, implementing it with code, and writing the paper.

**Han Li**<sup>†</sup>: Designing the architecture of ABTI module, implementing it with code, and writing the paper.

**Yuqing Wang**: Running ablation experiments of AQFormer and writing the experiment part.

**Hongji Zhu**: Running some of the ablation experiments of AQFormer and the ABTI module.

**Jiao Dai**<sup>\*</sup>: Participating in the discussions regularly, providing guidance, and writing the introduction part.

**Jizhong Han**: Participating in the discussions regularly and writing the method part.

**Wenge Rong**: Participating in the discussions regularly and writing the related work part.

**Si Liu**: Participating in the discussions regularly and writing the abstract part.

<sup>†</sup> means equal contributions and <sup>\*</sup> means corresponding author.

## References

- [Alwassel *et al.*, 2020] Humam Alwassel, Dhruv Mahajan, Bruno Korbar, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. *NeurIPS*, 2020.
- [Arandjelovic and Zisserman, 2017] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *ICCV*, 2017.
- [Arandjelovic and Zisserman, 2018] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *ECCV*, 2018.
- [Botach *et al.*, 2022] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multimodal transformers. In *CVPR*, 2022.
- [Carion *et al.*, 2020] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [Chen *et al.*, 2017] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [Cheng *et al.*, 2021] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *NeurIPS*, 2021.
- [Cheng *et al.*, 2022] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022.
- [Cho *et al.*, 2014] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- [Ding *et al.*, 2022] Zihan Ding, Tianrui Hui, Junshi Huang, Xiaoming Wei, Jizhong Han, and Si Liu. Language-bridged spatial-temporal interaction for referring video object segmentation. In *CVPR*, 2022.
- [Dosovitskiy *et al.*, 2020] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [Duke *et al.*, 2021] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *CVPR*, 2021.
- [Fayyaz *et al.*, 2016] Mohsen Fayyaz, Mohammad Hajizadeh Saffar, Mohammad Sabokrou, Mahmood Fathy, Reinhard Klette, and Fay Huang. Stfcn: spatio-temporal fcn for semantic video segmentation. *arXiv preprint arXiv:1608.05971*, 2016.
- [Fu *et al.*, 2019] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, 2019.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Hershey *et al.*, 2017] 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 *ICASSP*, 2017.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 1997.
- [Hu *et al.*, 2019] Di Hu, Feiping Nie, and Xuelong Li. Deep multimodal clustering for unsupervised audiovisual learning. In *CVPR*, 2019.
- [Lin *et al.*, 2014] 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 *ECCV*, 2014.
- [Long *et al.*, 2015] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [Milletari *et al.*, 2016] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *3DV*, 2016.
- [Nilsson and Sminchisescu, 2018] David Nilsson and Cristian Sminchisescu. Semantic video segmentation by gated recurrent flow propagation. In *CVPR*, 2018.



- [Oh *et al.*, 2019] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *ICCV*, 2019.
- [Perazzi *et al.*, 2016] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, pages 724–732, 2016.
- [Qian *et al.*, 2020] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. In *ECCV*, 2020.
- [Rouditchenko *et al.*, 2019] Andrew Rouditchenko, Hang Zhao, Chuang Gan, Josh McDermott, and Antonio Torralba. Self-supervised audio-visual co-segmentation. In *ICASSP*, 2019.
- [Russakovsky *et al.*, 2015] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 2015.
- [Seo *et al.*, 2020] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *ECCV*, 2020.
- [Strudel *et al.*, 2021] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *ICCV*, 2021.
- [Wang *et al.*, 2022] 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*, 2022.
- [Wu *et al.*, 2022a] Ho-Hsiang Wu, Prem Seetharaman, Kundan Kumar, and Juan Pablo Bello. Wav2clip: Learning robust audio representations from clip. In *ICASSP*, 2022.
- [Wu *et al.*, 2022b] Jiannan Wu, Yi Jiang, Peize Sun, Zehuan Yuan, and Ping Luo. Language as queries for referring video object segmentation. In *CVPR*, 2022.
- [Yang and Yang, 2021] Zongxin Yang and Yi Yang. Associating objects with transformers for video object segmentation. *NeurIPS*, 2021.
- [Yang and Yang, 2022] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. *NeurIPS*, 2022.
- [Yang *et al.*, 2020] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *ECCV*, 2020.
- [Yuan *et al.*, 2021] Yuhui Yuan, Lang Huang, Jianyuan Guo, Chao Zhang, Xilin Chen, and Jingdong Wang. Ocnet: Object context for semantic segmentation. *IJCV*, 2021.
- [Zhao *et al.*, 2019] Hang Zhao, Chuang Gan, Wei-Chiu Ma, and Antonio Torralba. The sound of motions. In *ICCV*, 2019.
- [Zhou *et al.*, 2022] 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 *ECCV*, 2022.
- [Zhu *et al.*, 2017] Xizhou Zhu, Yuwen Xiong, Jifeng Dai, Lu Yuan, and Yichen Wei. Deep feature flow for video recognition. In *CVPR*, 2017.
- [Zhu *et al.*, 2021] X Zhu, W Su, LW Lu, B Li, XG Wang, and Deformable DETR Dai J F. Deformable transformers for end-to-end object detection. In *ICLR*, 2021.