

Unveiling and Mitigating Bias in Audio Visual Segmentation

Anonymous Authors

ABSTRACT

Community researchers have developed a range of advanced audio-visual segmentation models aimed at improving the quality of sounding objects' masks. While masks created by these models may initially appear plausible, they occasionally exhibit anomalies with incorrect grounding logic. We attribute this to real-world inherent preferences and distributions as a simpler signal for learning than the complex audio-visual grounding, which leads to the disregard of important modality information. Generally, the anomalous phenomena are often complex and cannot be directly observed systematically. In this study, we made a pioneering effort with the proper synthetic data to categorize and analyze phenomena as two types "audio priming bias" and "visual prior" according to the source of anomalies. For audio priming bias, to enhance audio sensitivity to different intensities and semantics, a perception module specifically for audio perceives the latent semantic information and incorporates information into a limited set of queries, namely active queries. Moreover, the interaction mechanism related to such active queries in the transformer decoder is customized to adapt to the need for interaction regulating among audio semantics. For visual prior, multiple contrastive training strategies are explored to optimize the model by incorporating a biased branch, without even changing the structure of the model. During experiments, observation demonstrates the presence and the impact that has been produced by the biases of the existing model. Finally, through experimental evaluation of AVS benchmarks, we demonstrate the effectiveness of our methods in handling both types of biases, achieving competitive performance across all three subsets.

CCS CONCEPTS

• Computing methodologies → Scene understanding.

KEYWORDS

Audio-Visual Segmentation, Multimodal Bias, Multimodal Learning

1 INTRODUCTION

In the real world, audio and visual are two closely related modalities that provide fundamental perception in regular life, often applied in forms such as videos. One classic task within the field of audio-visual understanding is Audio-Visual Localization (AVL) [17, 36], which enables the unsupervised localization of sounding objects by utilizing sound as guidance. With the increasing demand for stronger perceptual capabilities in autonomous driving [37, 48], and

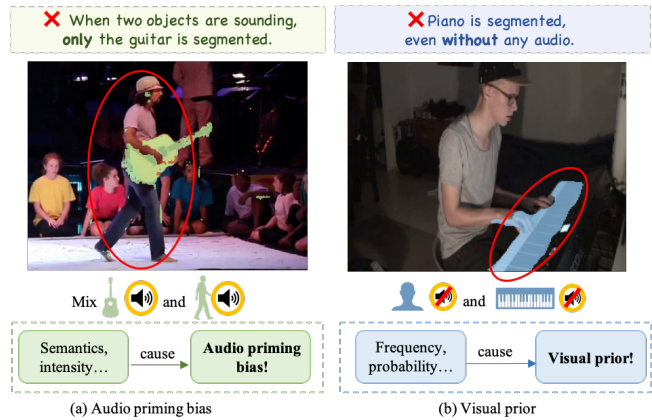


Figure 1: The illogical anomalies caused by biases. (a) Even if each audio yields a satisfactory mask separately, the dominance still exists when they overlap. (b) In training data, the piano generally sounds. During testing, regardless of the sound presented, the model still tends to prioritize the piano. These two phenomena of the illogical anomalies can be categorized into "audio priming bias", and "visual prior", which typically are simultaneously observed as a general impediment in the AVS model.

embodied intelligence [13, 31], there is an urgent to establish finer grounding beyond bounding boxes and heat maps between audio and visual elements. Therefore, a new task called Audio-Visual Segmentation (AVS) has emerged. It introduces a pixel-level, fine-grained scene understanding, bringing novel challenges to the field of audio-visual understanding. Furthermore, the performance of AVS showcases the current ability of machines to understand and integrate modalities in complex scenarios.

The existing works on AVS can be broadly categorized into fusion-based [18, 22, 25, 28, 46] and prompt-based methods [10, 29, 30, 38] according to Wang *et al.* [38]. The former primarily integrates audio and visual information at different stages to localize sounding objects, while the latter focuses on generating effective audio prompts and further finetuning the visual foundation model pre-trained on large segmentation datasets [20, 44]. Currently, prompt-based methods are the cutting-edge models in this field, demonstrating exceptional performance. However, researchers have observed illogical anomalies in these approaches [38, 43]. As illustrated in Fig. 1, even when the mask of the object seems plausible, two anomalous but interesting phenomena can be observed. 1) The segmentation results after overlaying audio do not align with the superimposition of masks guided by separate audios, even when each audio yields a satisfactory mask. Instead, dominance occurs. 2) Regardless of whether the piano produces sound, the model usually tends to segment the piano. Phenomena like Fig. 1 have a wide impact on the current model, which makes analyzing and addressing these phenomena an urgent work to improve the grounding behavior.

Permission to make digital or hard copies of all or part of this work for personal or professional use, is granted by ACM Publishing Department for non-commercial use, provided that the copyright notice and the full citation on the first page are preserved. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM MM, 2024, Melbourne, Australia
© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
<https://doi.org/10.1145/nmmmmmm.nmmmmmm>

Generally, the phenomena depicted on the left and right sides of Fig. 1 are often coupled and cannot be directly observed individually. We manage to control the influencing factor with artificial data, including volume, etc., and reasonably categorize phenomena into two types of biases: “audio priming bias” and “visual prior”. Specifically, “**audio priming bias**” refers to the phenomenon that the model tends to focus on audio salient content but not whole content. This bias is characterized by the model’s insensitivity to audio of specific intensity or semantics and consequently the dominance of certain semantic audio in multi-source scenarios. As for “**visual prior**”, it refers to the phenomenon that the model may directly segment the common-sounding objects. The most common observed form of such bias is that regardless of the audio guidance employed, the model consistently segments the whole or part of certain objects. It is evident that both of these biases significantly impact the behavior of the current AVS models.

To provide the community with a comprehensive solution to address these biases, we adopt the divide-and-conquer thought. For audio priming bias, we first introduce **semantic-aware active queries**. These active queries contain rich latent semantic information gathered by the perception module, enhancing the sensitivity to audio cues. Furthermore, we employ a customized **interaction mechanism** in the transformer decoder specially designed to enhance the active queries and suppress the dormant ones for better collaboration of specific audio semantics. As for visual priors, we explore three types of **contrastive debias strategies** without changing the structure of the model. Our findings indicate that the soft and gradual debias strategy based on uncertainty yields superior results across various AVS methods including ours.

Finally, through experimental evaluation of AVS benchmarks, we demonstrate the effectiveness of our methods in handling both types of biases, achieving highly competitive performance across all three subsets. In summary, our contributions are threefold:

- We make a pioneering attempt to categorize the complex phenomena into two types of bias: “audio priming bias” and “visual prior” with necessary observation and analysis.
- For audio priming bias, we propose semantic-aware active queries and customized interaction mechanisms to improve the sensitivity and cooperation of complex audio scenarios. Then, for visual prior, the novel debias strategies are utilized to contrastively reorganize the distribution of logits.
- Experimental results reveal that the proposed method and strategy effectively mitigate biases while achieving versatility and comparable performance.

2 RELATED WORKS

2.1 Audio-Visual Segmentation

Before the emergence of AVS, traditional AVL tasks [17, 32, 36] relied on unsupervised learning that used bounding boxes or coarse heat maps for predicting the positions of sounding objects. However, for applications such as autonomous driving [37, 48], and embodied intelligence [13, 31], a more precisely multimodal grounding task to localize sounding objects at a pixel level is required. This is where AVS comes into the stage. The existing AVS approaches can be broadly categorized into fusion-based [18, 22, 25, 28, 43, 46] and prompt-based [10, 29, 38, 41]. The pioneering fusion-based

work [46] employs a multi-stage strategy to integrate audio with multi-scale visual features. Building upon this, CATR [22] further proposes a segmentation paradigm that combines both temporal and fusion information. Taking the thought of bidirectional generation, Hao *et al.* [14] achieve a further boosting on the performance. In contrast, prompt-based methods like AVSegFormer [10] and GAVS [38] directly decode fused features using projected audio queries. Notably, GAVS leverages rich visual information to achieve a generalizable model in zero-shot and few-shot scenarios. Moreover, taking a holistic perspective, Yan *et al.* [41] present a general framework for AVS and Ref-VOS [39] (Referring Video Object Segmentation), which both are cross-model guided video segmentation. However, previous AVS studies have primarily focused on achieving better multimodal fusion and plausible masks from the framework perspective, without delving into the challenges underlying the establishment of reasonable grounding.

2.2 Bias in Cross-Model Guided Tasks

Due to the widespread interest in the Visual Question Answering (VQA) task [3, 4], researchers have tackled two types of bias, namely “visual priming bias” and “language prior”, as discussed in previous works [33, 42]. “Visual priming bias” was initially identified by Antol *et al.* [4], where models tend to prioritize visually salient content while disregarding the other context [12, 42]. Subsequently, the issue of “language prior” has gained attention from researchers [2, 33, 42], referring to the tendency of models to generate statistically common answers (e.g. the problem of green and yellow banana). The approaches for addressing the aforementioned bias can be broadly categorized into non-extra-based methods [5, 8, 33], and extra-based methods, involving additional human supervision [35, 40]. Different from extra-based methods, non-extra-based methods introduce a branch of the question-only model and conduct joint training with the target VQA model by specific learning strategies. Researchers have proposed several customized learning strategies including sharing the same question encoder [33], blocking the backpropagation of the question-only model [5], pretraining question-only model [8]. Moreover, it has been discovered recently that there is also a certain similarity between bias in VQA and hallucination [21, 47] in the popular Multimodal Large Language Model (MLLM).

However, the bias in AVS has never been explored by researchers before. Here, we hold the holistic opinion that VQA and AVS, as cross-model guided tasks, share certain similarities. VQA involves obtaining answers to questions based on image cues, while AVS involves obtaining masks from images based on audio cues. Therefore, in simple terms, VQA can be understood as a symmetrical problem to AVS in terms of modalities. So, the anomalous phenomena observed in AVS can potentially be explained by the bias theory present in VQA.

3 PROBLEM AND ANALYSIS

Based on the same thought of bias categories in VQA [42], we similarly classify the phenomena of anomalies into two types, “audio priming bias”, and “visual prior” in Fig. 1. The segmentation failure cases observed in the previous model often tend to be a combination of both biases.

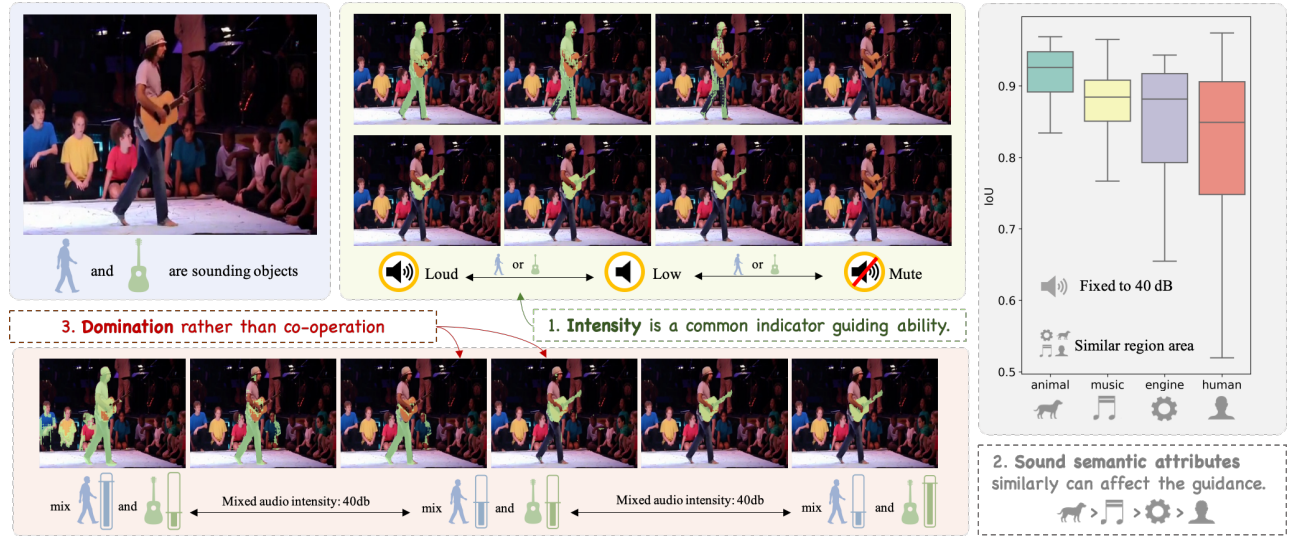


Figure 2: The illustration of audio priming bias. We have discovered that higher audio intensity results in stronger guiding capabilities in the green block. Also, the audio with distinct semantic attributes is easier to learn and possesses stronger guiding capabilities in the grey block. For instance, music has a greater guiding capability compared to human sound. Consequently, the diverse guiding capabilities cause the phenomenon of dominance in the red block.

3.1 Audio Priming Bias

The phenomenon that the model tends to focus on audio salient content but not whole content is called “audio priming bias”. Firstly, through observation experiments involving manual intervention in audio (detailed in the appendix), we have observed the phenomenon of audio priming bias as illustrated in Fig. 2. Hence, the following observations can be discerned. 1) Audio with different intensities demonstrates varying guiding capability, as shown in Fig. 2 green block. 2) When controlling other variables including volume, we can observe a clear variance of the guiding capability by different semantics through the box plot. 3) In cases where multiple audios are simultaneously present, the overlaying of audio does not always lead to separate related masks being superimposed. Instead, the dominant of the audio appears as shown in Fig. 2 red block. Therefore, we contend that the primary determinants of audio guiding abilities are the **intensity** and **semantic attributes** of the audio. Consequently, the diverse guiding capabilities give rise to the phenomenon of dominance. Therefore, a well-designed mechanism is demanded to alleviate audio priming bias.

3.2 Visual Prior

The phenomenon that the model may directly segment the common-sounding objects is called “visual prior”. Figure 4 reveals that, regardless of the audio provided, the AVS model consistently segments the partial or entire region of the particular object. Since the piano generally appears with a high sounding probability in the training data, the model tends to segment once sees the piano. Statistically, the **occurrence frequency** in [45] and the **sounding probability** in Fig. 4 between different semantics are imbalanced in the dataset. According to previous works [21, 23, 33], such preference and distribution provide strong prior information and make the model inclined to obtain statistically plausible results, rather than achieving the desired challenging grounding behavior. Moreover,

due to the imbalance of inherent preference and distribution in the real world, addressing visual prior becomes more significant.

4 METHODS

The overall pipeline for the AVS task is normally constructed by the encoder-decoder structure in Fig. 3 involving an image encoder, a pixel decoder, and a transformer decoder. Generally, the multimodal information is fused and perceived in the transformer decoder [7, 10, 19, 22]. Therefore, the mitigation of **audio priming bias** requires an enhancing mechanism in the transformer decoder, while the mitigation of **visual prior** requires distribution reorganization after acquiring the logits. Methodologically, on one hand, to enhance the audio sensitivity and cooperation of different intensity and semantic attributes, we introduce the semantic-aware active queries utilizing a perception module and interaction enhancement mechanism. On the other hand, multiple training strategies are explored to contrastively optimize the debias model and reorganize the logits without modifying the structure.

4.1 Preliminaries

Given an audio-visual video sequence, we first divide it into t non-overlapping audio and visual segment pairs with a length of 1s.

Visual. After the division, we extract visual representations $F_V \in \mathbb{R}^{t \times d_V \times H \times W}$ using a pre-trained Swin-base model [9]. The visual encoder is frozen and only tuned with two-layer MLP adapters.

Audio. Similarly, audio representations $F_A \in \mathbb{R}^{t \times d_A}$ is encoded from VGGish [11, 15], where t represents the duration of the audio in seconds and matches the number of video frames. The audio representations are extracted in advance by the frozen encoder.

Task. Given a dataset $\mathcal{D} = \{F_V^i, F_A^i, M_{gt}^i\}_{i=1}^n$ consisting of triplets of images $F_V^i \in F_V$, audio $F_A^i \in F_A$ and masks $M_{gt}^i \in M_{gt}$, the AVS task is to learn the mapping $\mathcal{G}(F_V, F_A) \rightarrow M_{gt}$.




Figure 3: In conventional semantic segmentation, the learnable queries are solely responsible for generating regions. However, in an ideal AVS model with a similar structure, the learnable queries need to not only generate regions but also perceive and regulate the interaction between audio semantics. So, this mechanism is designed to enhance the understanding of latent semantics and improve the interaction to only occur between semantic-aware active queries.

4.2 Semantic-Aware Active Queries

In transformer decoder of conventional semantic segmentation, each learnable query ($\in \mathbb{R}^{1 \times d_A}$ among initialized learnable queries $q \in \mathbb{R}^{t \times N \times d_A}$, where N is the number of learnable queries in 1s interval) functions like a region proposal network [34] and has the ability to generate mask proposals. These queries are simply inherited by repeatedly adding audio features F_A , in the AVS model [10, 22, 38] as shown in the upper block in Fig. 2. However, in an ideal AVS model, the learnable queries need to not only generate regions but also perceive and regulate the interactions between audio semantics adaptively.

To achieve the basic perception of audio semantics, we first introduce the following perception module aiming to enhance sensitivity to audio of various intensities and semantics. The logical and straightforward approach is to employ a separately trained projection $\Phi(\cdot)$ module with an optimization objective for audio alone to perform a coarse perception of latent audio semantics and to finally project to a latent semantic mark $\Phi(F_A) \in \mathbb{R}^{t \times N}$ with padding of 0.

$$\Phi(F_A) = \{1, x_1, x_2, \dots\}, \quad (1)$$

$$x_i = \begin{cases} 1 & \text{if } F_A \in i^{th} \text{ class or cluster,} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Here, we use K-means, Gaussian Mixture Models (GMMs), or BEATS [6] to perform projection in both unsupervised and supervised ways. So each audio segment will be assigned with $1 \leq C \leq 8$ classes or $C = 1$ clusters. The overall goal of a latent semantic mark is a coarse filter to have a coarse and basic perception of the audio semantics. The items encoded by 1 in $\Phi(F_A)$ are termed *active* marks, whereas those encoded by 0 in $\Phi(F_A)$ are termed *dormant* marks. The active or dormant states are relative and vary from samples. Afterward, each audio segment will receive the latent semantic marks $\Phi(F_A) \in \mathbb{R}^{t \times N}$ based on the projection mentioned above. Then the semantic-enhanced audio features are obtained by element-wise product \odot with broadcast

$$F'_A = \text{repeat}(F_A, N) \odot \Phi(F_A), \quad (3)$$

where $F'_A \in \mathbb{R}^{t \times N \times d_A}$ is semantic-enhanced audio features. Furthermore, we consistently encode the 0^{th} mark out of N as an active mark shown in Eq. 1 to capture global audio semantic information.

Building upon the semantic-enhanced audio features, we carry out self-attention $\text{SA}(\cdot)$ with the residual of initialized learnable queries q and obtain the semantic-aware queries q' as

$$q' = \text{SA}(F'_A, F'_A, F'_A) + q = \text{softmax}\left(F'_A F'_A{}^T\right) F'_A + q, \quad (4)$$

where $q' \in \mathbb{R}^{t \times N \times d_A}$ is the semantic-aware queries. Here, to enhance conciseness, the equations in this paper for attention have excluded projection and scaling factors. The semantic-aware queries obtained are utilized as the learnable query inputs for the transformer decoder layers. Similarly, the queries in the same position corresponding to latent semantic active marks are *active queries* q'_a , while those corresponding to dormant marks are *dormant queries* q'_d , where $\{q'\} = \{q'_a\} \cup \{q'_d\}$. Moreover, the adequate number of active queries q'_a per segment is ensured in multiple ways and please refer to the appendix for more details.

4.3 Enhancing Interaction within Active Queries

After introducing active and dormant queries, we observed that regions proposed by dormant queries are sometimes unnecessarily optimized and backpropagated. It is clearly detrimental since dormant queries do not contain adequate audio semantic information. The widely used transformer decoder layer in per-mask segmentation consists of three modules to process learnable query features $q \in \mathbb{R}^{N \times t \times d_A}$ in the following order: a cross-attention $\text{CA}(\cdot)$, a self-attention module $\text{SA}(\cdot)$, and a feed-forward network $\text{FFN}(\cdot)$. For the original transformer decoder, the interactions occur between all learnable query features, and the same query can propose regions of different semantics. Original cross-attention and self-attention in l^{th} layer can be denoted as

$$q_{(l)}^{ca} = \text{CA}(q_{(l-1)}, F_{V(l)}, F_{V(l)}) + q_{(l-1)} \\ = \text{softmax}(\mathbf{m}_{(l-1)} + q_{(l-1)} F_{V(l)}) F_{V(l)} + q_{(l-1)}, \quad (5)$$

$$q_{(l)}^{sa} = \text{SA}(q_{(l)}^{ca}, q_{(l)}^{ca}, q_{(l)}^{ca}) + q_{(l)}^{ca} \\ = \text{softmax}(q_{(l)}^{ca} q_{(l)}^{ca T}) + q_{(l)}^{ca}, \quad (6)$$

where l is the layer index, $q_{(l)}$ refers to the query features corresponding to the learnable queries q . $q_{(l)}^{ca}$ and $q_{(l)}^{sa}$ is the temporary

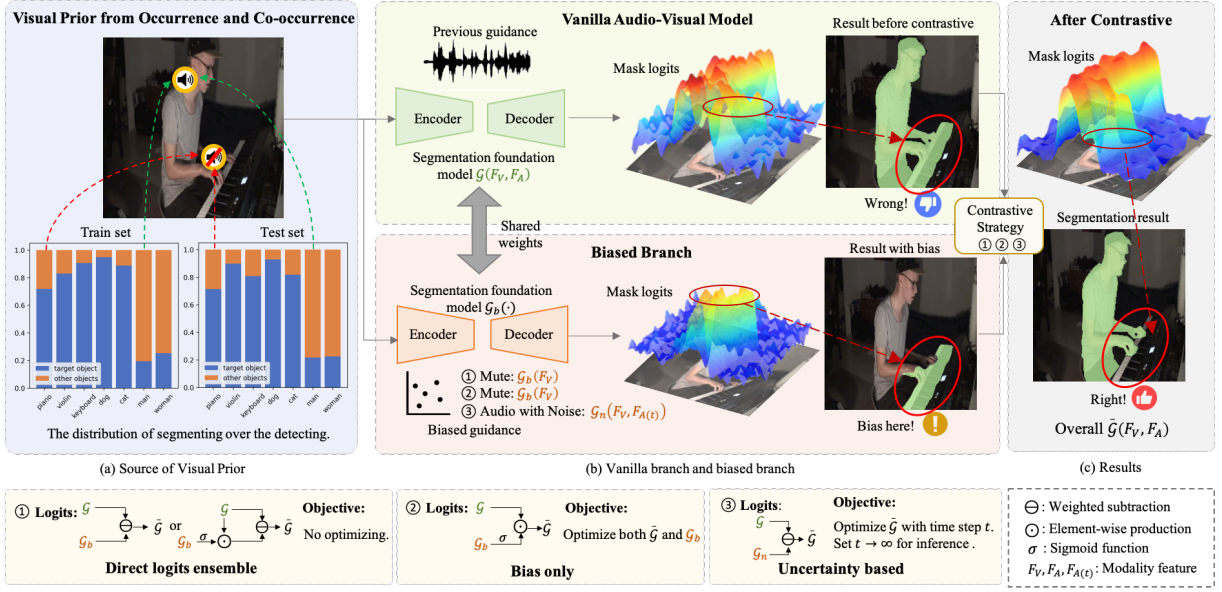


Figure 4: Illustration of visual prior. (a) The model always tends to learn statistically plausible results, rather than achieve the harder desired grounding behavior. Note: the bar chart uses blue to represent the proportion of the object being present in the image and emitting sound, while the orange color represents the proportion of the object being present but not emitting sound. As an example, since the piano generally appears with a high sounding probability in the training data, the model will segment once sees the piano. (b) To deal with the visual prior, we introduce debias strategies through the idea of contrasting the audio-visual model with the biased branch. (c) The ideal result is the mask without visual prior.

state after the cross-attention $CA(\cdot)$ and after self-attention $SA(\cdot)$ respectively. Note, $F_{V(l)}$ are the image features under transformation, $\mathbf{m}(l)$ is the masked attention introduced by Chen *et al.* [7].

However, we argue that audio interactions should only occur between active queries related to perceived latent semantics. This ensures that a specific set of active queries should only propose the region of the same semantics.

$$q_a^{ca}(l) = CA(q_a^{l-1}, F_{V(l)}, F_{V(l)}) + q_a^{l-1}, \quad (7)$$

$$q_a^{sa}(l) = SA(q_a^{ca}(l), q_a^{ca}(l), q_a^{ca}(l)) + q_a^{ca}(l), \quad (8)$$

where only active query features q_a^{l-1} are updated to $q_a^{ca}(l)$ by $FFN(q_a^{ca}(l))$, and dormant query features q_d^{l-1} , are masked during backpropagation. Finally, based on the $q_a^{ca}(l)$ in each layer, the mask logits $\in \mathbb{R}^{T \times N \times H \times W}$ and class logits $\in \mathbb{R}^{T \times N \times \mathcal{K}}$ are obtained, where \mathcal{K} is the number classes of semantics ($\mathcal{K}=1$ for binary).

Following each transformer decoder layer, query matching aims to determine which of the predicted regions fits the referred objects. Therefore, we match the predicted regions from the active queries and the ground truth regions by minimizing the matching cost

$$(\mathcal{M}, \mathcal{M}_{gt}) = \arg \min_{M \in \mathcal{M}_{q_a}} C_{\text{match}}(\mathcal{M}_{q_a}, \mathcal{M}_{gt}), \quad (9)$$

where \mathcal{M}_{q_a} are the regions proposed by q_a , and \mathcal{M}_{gt} is the ground truth of original image. \mathcal{M}_{q_a} is calculated by multiplying mask logits and class logits, so the cost C_{match} evaluates the predicted region in both mask and semantic level. In practice, we separately compute the costs of mask and class logits and jointly optimize the summation as C_{match} . Then, the losses \mathcal{L} are calculated in every

intermediate transformer decoder layer, and only the one in the final layer is given a higher weight and used for inference.

Based on the method above, interactions related to active queries are retained and enhanced, while interactions of dormant queries are suppressed with a reduction of the computational workload. Although this method slightly reduces the number of proposed valid regions \mathcal{M}_{q_a} , this trade-off is proven advantageous for better cooperation between audio semantics in further experiments.

4.4 Contrastive Debias Strategy

To recognize bias and then avoid the biased region appearing along with the referred region, it is a natural idea to introduce a new biased branch and conduct joint training to the current AVS model without modifying their structures. We first briefly bring up two simple and direct methods and carry out analysis and experiment for comparison. Subsequently, a well-designed uncertainty-based strategy is proposed with competitive performance and versatility. To maintain brevity, operations on logits in Sec. 4.4 are applied to both mask logits and class logits separately.

To begin with, the simplest approach to reorganize the distribution is the **direct logits ensemble strategy** of separately trained biased model $\mathcal{G}_b(F_V) \rightarrow \mathcal{M}_{gt}$ and vanilla AVS model $\mathcal{G}(F_V, F_A) \rightarrow \mathcal{M}_{gt}$, and subsequently merge their output logits using simple operations such as weighted subtraction and multiplication. As the second strategy, the **bias-only strategy** requires a new objective of training by involving a bias-only branch with the mapping of $\mathcal{G}_b(F_V) \rightarrow \mathcal{M}_{gt}$ and jointly optimize the overall model with a bias-only branch. To obtain the debias logits, we simply

compute an element-wise product \odot of the logits of \mathcal{G} and \mathcal{G}_b

$$\text{logits}(\tilde{\mathcal{G}}(F_V, F_A)) = \text{logits}(\mathcal{G}(F_V, F_A)) \odot \sigma(\text{logits}(\mathcal{G}_b(F_V))), \quad (10)$$

where $\tilde{\mathcal{G}}$ is the debias projection of the triplets \mathcal{D} and σ is the sigmoid function. The final optimization goal can be calculated from the logits as

$$\mathcal{L}_{\text{bias}}(\mathcal{G}, F_V, F_A, \mathcal{G}_b) = \mathcal{L}(\tilde{\mathcal{G}}, F_V, F_A) + \mathcal{L}(\mathcal{G}_b, F_V), \quad (11)$$

Based on the updated logits, the region of the man in Fig. 4 will be given fewer logits and yield a greater loss than the original.

Although plausibly reorganizing the logits distribution, these two methods demonstrate fluctuations in experiments and do not consistently improve performance in popular methods. These issues can be attributed to the hard and direct introduction of the biased branch, which forces the network to learn from the large distribution gap, resulting in fluctuations in gradient direction and magnitude. To address this, we propose a soft and gradual strategy by introducing audio uncertainty through Gaussian noise.

Our proposed **uncertainty-based strategy** provides a gradual and soft way to estimate the biased output distribution. It does not require multitask-like loss but involves contrast \mathcal{G}_n output distributions derived from image and distorted audio inputs. We first follow the idea of the forward diffusion process $r(F_{A(T)} | F_{A(0)})$ in image generation [16] to construct the distorted audio

$$\begin{aligned} r(F_{A(\tau)} | F_{A(\tau-1)}) &= \mathcal{N}(F_{A(\tau)}; \sqrt{1-\beta}F_{A(\tau-1)}, \beta\mathbf{I}), \\ r(F_{A(T)} | F_{A(0)}) &= \prod_{\tau=1}^T r(F_{A(\tau)} | F_{A(\tau-1)}), \end{aligned} \quad (12)$$

where $F_{A(0)}$ denotes the original audio features without any noise and I refers to an identity matrix. We incrementally add a small amount of Gaussian noise for T steps. As step τ goes larger, the amount of noise added in each step is controlled by β . Finally, if $T \rightarrow \infty$, $F_{A(T)}$ will be completely distorted. Then, the final logits are updated as

$$\begin{aligned} \text{logits}(\tilde{\mathcal{G}}(F_V, F_A)) &= (\alpha + 1)\text{logits}(\mathcal{G}(F_V, F_A)) \\ &\quad - \alpha \text{logits}(\mathcal{G}_n(F_V, F_{A(T)})), \end{aligned} \quad (13)$$

where larger α values indicate a stronger amplification of differences between the two distributions ($\alpha = 0$ reduces to regular logits). During inference, $T \rightarrow \infty$, so $F_{A(T)}$ only contains Gaussian noise for the biased branch. Based on such logits, the overall loss

$$\mathcal{L}_{\text{uncertainty}} = \mathcal{L}(\tilde{\mathcal{G}}, F_V, F_A), \quad (14)$$

is calculated by updated logits. Finally, the loss and mask result will be calculated from the debias logits $\text{logits}(\tilde{\mathcal{G}}(F_V, F_A))$. In summary, replacing the hard bias-only branch with a gradually distorted audio branch provides a soft and learnable transition for the model to reorganize the distribution caused by visual prior.

4.5 Loss and Inference

Considering both masks and semantics, the overall loss \mathcal{L} mentioned in Sec.4.4 is obtained within every intermediate transformer decoder layer by mask loss $\mathcal{L}_{\text{mask}}$ [24] and class loss \mathcal{L}_{cls} based

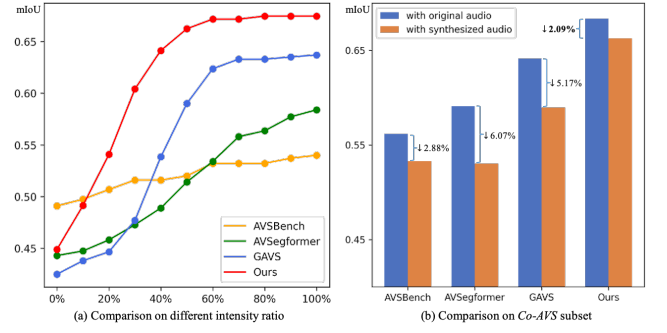


Figure 5: (a) The performance comparison of different methods on V1M under certain intensity conditions. Our method brings more sensitivity to low-intensity scenarios. (b) The performance comparison of different methods on our Co-AVS subsets. Our method has advantages in both performance and robustness.

on the updated mask logits and class logits by debias strategy as

$$\mathcal{L} = \mathcal{L}_{\text{mask}} + \lambda_{\text{cls}} \cdot \mathcal{L}_{\text{cls}}, \quad (15)$$

$$\mathcal{L}_{\text{mask}} = \lambda_{\text{focal}} \cdot \mathcal{L}_{\text{focal}} + \lambda_{\text{dice}} \cdot \mathcal{L}_{\text{dice}}, \quad (16)$$

where $\mathcal{L}_{\text{mask}}$ is implemented by a combination of the dice loss $\mathcal{L}_{\text{dice}}$ [29] and the focal loss $\mathcal{L}_{\text{focal}}$ functions, and \mathcal{L}_{cls} is implemented by the cross-entropy loss. After adopting the active queries, both losses $\mathcal{L}_{\text{mask}}$, \mathcal{L}_{cls} are calculated and backpropagated within the matched pairs of masks \mathcal{M} from active queries and ground truth \mathcal{M}_{gt} . During inference, the overall logits $\in \mathbb{R}^{(K+1) \times H \times W}$ are obtained by the multiplication of mask logits and class logits in the last layer.

To address the shortcut effect resulting from easily learned regions proposed by specific queries, we drop the relatively well-learned pairs $(\mathcal{M}, \mathcal{M}_{gt})$ under the probability of p in the calculation of \mathcal{L} and empirical results show slight improvements.

5 EXPERIMENTS

5.1 Implementation Details

Dataset. Our proposed method is evaluated on the AVS Benchmarks [46], which contains three subsets. Firstly, the single-source subset (V1S) contains 4932 videos over 23 categories. In this subset, only the first sampled frame is annotated. Secondly, the multi-source subset (V1M) contains 424 videos that include two or more sound sources and all sounding objects are visible in the frames. Finally, the AVSBench-semantic (AVSS) subset, an extension of V1S and V1M, contains 12,356 videos of 10s, differing from the 5s videos in V1S and V1M. Additionally, to validate the cooperation performance in complex audio scenarios, a small multi-source test set of AVS called Co-AVS is created, containing 450 videos of 5s. Original audio clips are replaced with synthesized audio clips of 6 different levels intensities of each semantics, like the red block in Fig. 2.

Setting. We conduct training and evaluation using the VGGish backbone pretrained on Youtube-8M [1] and Swin-base Transformer backbone pretrained on semantic-ADE20K [44] by Mask2Former [7]. The parameters λ_{focal} , λ_{dice} , λ_{cls} in the loss are set to 5, 5, 2, 1. The AdamW optimizer is adopted with a learning rate of 1e-4 for the visual encoder adapters and 1e-3 for other learnable parameters,

Table 1: Comparison of performance on all three subsets of AVS-Benchmarks. The bold score is utilized to represent optimal results. “-” refers to the methods that do not support the corresponding subsets. Except for specific references, the “strategy” mentioned in the comparison refers to the uncertainty-based strategy that achieves the best outcomes.

Method	Audio-backbone	Visual-backbone	V1S		V1M		AVSS	
			mIoU(%)	F-score	mIoU(%)	F-score	mIoU(%)	F-score
AVSBench [46]	VGGish	PVT-v2	78.70	0.879	54.00	0.645	29.77	0.352
AVSC [26]	VGGish	PVT-v2	81.29	0.886	59.50	0.657	-	-
AVS-BG [14]	VGGish	PVT-v2	81.71	0.904	55.10	0.668	-	-
AQFormer [18]	VGGish	PVT-v2	81.60	0.894	61.10	0.721	-	-
AuTR [28]	VGGish	Swin-Base	80.40	0.891	56.20	0.672	-	-
BAVS [27]	BEATs	Swin-Base	82.68	0.898	59.63	0.659	-	-
SAMA [29]	VGGish	ViT-Huge	81.53	0.886	63.14	0.691	-	-
† CATR [22]	VGGish	PVT-v2	81.40	0.896	59.00	0.700	36.66	0.420
AV-SAM [30]	ResNet18	ViT-Base	40.47	0.566	-	-	-	-
AVSBG [14]	VGGish	PVT-v2	81.71	0.904	55.10	0.668	-	-
GAVS [38]	VGGish	ViT-Base	80.06	0.902	63.70	0.774	-	-
AVSegFormer [10]	VGGish	PVT-v2	82.06	0.899	58.36	0.693	32.80	0.385
MUTR [41]	VGGish	Video-Swin-Base	81.60	0.897	64.00	0.735	-	-
Ours (w/o. strategy)	VGGish	Swin-base	82.92	0.928	66.12	0.792	41.93	0.476
Ours (w/. strategy)	VGGish	Swin-base	83.31	0.930	67.22	0.808	44.42	0.499

† CATR: To make a fair comparison, the results of CATR here are without supplemented annotation of the training set.

Table 2: Ablation study of the methods and strategies proposed in this study on V1M.

Ablation	mIoU(%)	F-score
Ours	67.22	0.808
w/o. Contrastive debias strategy	66.12	0.792
w/o. Dropping dominant queries	65.95	0.790
w/o. Interaction enhancement	62.57	0.744
w/o. Active query of 0^{th}	62.36	0.743
w/o. Active queries	62.16	0.734

with 60 training epochs. N and k in Fig. 3 is set to 100 and 9. Drop possibility p in Sec. 4.5 and α in Eq. 13 are set to 0.4 and 1 without elaborate searching. T related to the uncertainty level in Eq.12 are set to 1000.

Metrics. To conduct a comprehensive evaluation of our model, we carry out tests using mean Intersection over Union (mIoU) and F-score as the performance metrics [10, 38, 46].

5.2 Results Comparison

When compared to methods that do not incorporate the designed debias method, our method, which leverages methods for both biases, demonstrates strong competitiveness. As illustrated in Tab. 1, our method achieves comparable results across all V1S, V1M, and AVSS subsets, showcasing performance improvements of up to 0.63%, 3.22%, 7.76% in mIoU, respectively. Even when considering slight variations in the backbones used by different methods, our approach consistently outperforms others. The limited improvement observed in V1S can primarily be attributed to the constrained manifestation of biases stemming from the singular and simplistic nature of the data. However, we still believe that this does not undermine our ability to demonstrate the effectiveness of our bias handling techniques.

Furthermore, both quantitative performance in Fig. 5 and qualitative analysis in Fig. 6 indicate that our model successfully alleviates

Table 3: The analysis of projection strategy to obtain semantic-aware queries by class or cluster. The ablation of $topK$ can also be regarded as the ablation of a fixed number of active queries.

method	filtering	mIoU(%)	F-score
Classification	$\theta = 0.8$	64.55	0.757
	$\theta = 0.6$	66.36	0.783
	$\theta = 0.4$	67.22	0.808
	$\theta = 0.2$	64.52	0.740
	$topK = 1$	61.79	0.695
	$topK = 3$	64.49	0.742
Clustering	$topK = 5$	63.81	0.719
	Kmeans	64.18	0.746
	GMMs	64.57	0.742
None	None	63.59	0.702

both biases mentioned before and our predicted masks exhibit superior quality. Furthermore, due to the increased complexity of synthesized audio scenarios in Co-AVS subset, achieving a lower performance drop on synthesized audio necessitates better semantic collaboration. As shown in Fig.5, our demonstrates stronger robustness on complex audio scenarios.

5.3 Ablation Study

Overall ablation. We conduct an ablation experiment on the module and the strategy designed for biases. In this experiment, we progressively remove the modifications of each module in the set of V1M. Through this process, we aimed to assess the impact of these modifications on the segmentation results. As demonstrated in Tab. 2, the experiments indicate that the methods we designed for audio priming bias and visual prior respectively achieved an increase of 3.96% and 1.10% in mIoU. Different modules have all achieved an increase in both mIoU and F-score, reaching the design objectives. The ablation study further indicates that the inclusion of active queries alone does not directly result in a significant increase in

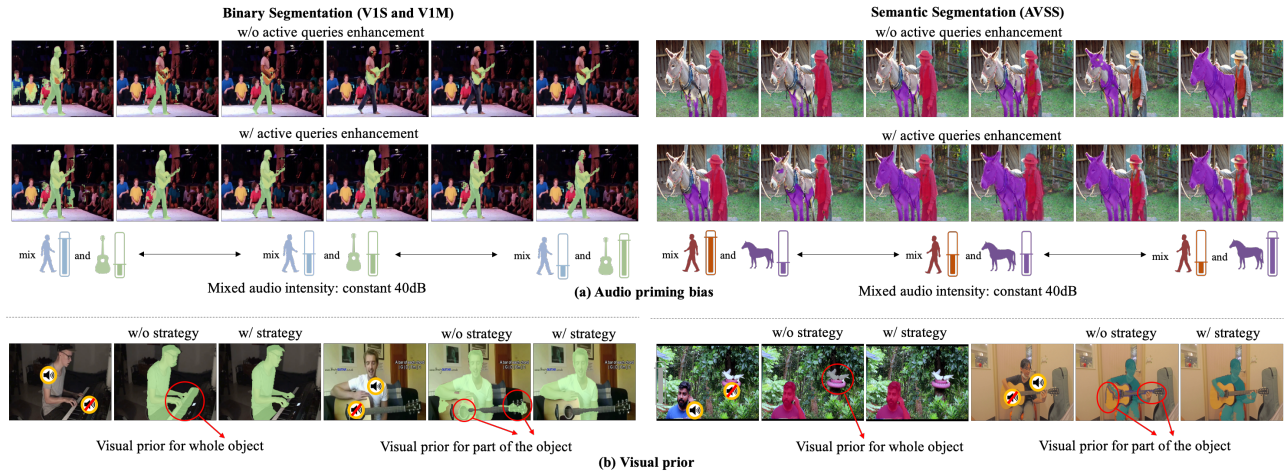


Figure 6: Our model successfully alleviates both biases and our predicted masks exhibit superior quality than before.

Table 4: Ablation analysis and the versatility experiments of image-biased strategy on V1M. “ \uparrow ” signifies a positive effect achieved by employing the contrastive debias strategy compared to the vanilla method, while “ \downarrow ” indicates a negative effect.

Method	Backbone	Vanilla		Logits ensemble		Bias only		Uncertainty based	
		mIoU (%)	Fscore	mIoU (%)	Fscore	mIoU (%)	Fscore	mIoU (%)	Fscore
AVSbench [46]	PVT-v2	54.00	0.645	54.12 \uparrow	0.659 \uparrow	53.43 \downarrow	0.640 \downarrow	54.81 \uparrow	0.682 \uparrow
AVSegformer [10]	PVT-v2	58.36	0.693	56.85 \downarrow	0.662 \downarrow	58.86 \uparrow	0.687 \downarrow	59.63 \uparrow	0.718 \uparrow
GAVS [38]	Swin-base	63.70	0.766	51.21 \downarrow	0.738 \downarrow	64.11 \uparrow	0.788 \uparrow	64.80 \uparrow	0.784 \uparrow
Ours	Swin-base	66.12	0.792	64.42 \downarrow	0.795 \uparrow	66.50 \uparrow	0.797 \uparrow	67.22\uparrow	0.808\uparrow

performance. It necessitates the implementation of enhancement of interaction, ultimately leading to a combined improvement of 3.76%. This finding emphasizes the importance of enhancing interaction in conjunction with active queries, as they aid the transformer encoder in establishing audio-visual grounding.

Ablation on projection of Equation 2. The objective of projection is to map each audio clip to several classes or a cluster with latent semantic information. Therefore, we naturally conducted experiments using both supervised and unsupervised approaches. In the supervised experiments, we employed Beats [6] for supervised multi-class classification and used threshold θ or topK strategies to roughly filter the samples with abstract semantic information in Tab. 3. The results showed that the different settings in the multi-classification do give an improvement of mIoU on V1M up to 67.22%. In the unsupervised experiments, we utilized the K-means and Gaussian Mixture Model (GMM) with Expectation Maximization for clustering, with the same number of clusters. Moreover, an adequate number of active queries per sample is ensured here for fairness. The results demonstrated that, regardless of supervised or unsupervised approaches, introducing latent semantic information along with active queries consistently improved performance, with a maximum improvement of 3.63% in Tab. 3.

Mask of single active queries. Besides the quantitative results in Tab. 2, it is qualitatively found that the masks generated by each active query are consistent with their respective semantics, which further validates the effectiveness of the semantic-aware active queries in audio perception. Detailed qualitative results of cases are provided in the appendix.

5.4 Versatility of Contrastive Debias Strategies

As general strategies, all three contrastive debias strategies are able to reorganize the distribution of logits. To validate the versatility, we not only applied the debias strategies in our framework but also explored its application in other popular models, as shown in Tab. 4. The performance of the logits ensemble and bias-only strategy fluctuates, achieving performance improvements only in certain methods. When considering all four methods, the uncertainty-based approach emerges as a superior solution due to its general improvement. Within our framework, it attained mIoU of 67.22% and exhibited commendable training stability.

6 CONCLUSION

This study presents a novel and systematic exploration of the bias issues existing in the previous AVS model. By defining and analyzing these biases as “audio priming bias” and “visual prior”, we propose targeted solutions for each bias. For audio priming bias, we introduce semantic-aware active queries that enhance the audio sensitivity of intensity and semantics with a designed perception module. Then further interaction of active queries is forced to collaborate on semantic understanding. For the visual prior, we employ a contrastive debias strategy that improves the overall performance of the model without modifying its structure. The meticulously designed methods and comprehensive experiment across all three subsets, demonstrate the superiority and versatility of our approach. In summary, our work introduces a novel perspective within the prevailing AVS network, opening up new possibilities for the future development of cross-model guided segmentation.

REFERENCES

- [1] Sami Abu-El-Hajja, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675* (2016).
- [2] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. Don't just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4971–4980.
- [3] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6077–6086.
- [4] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*. 2425–2433.
- [5] Remi Cadene, Corentin Dancette, Matthieu Cord, Devi Parikh, et al. 2019. Rubi: Reducing unimodal biases for visual question answering. *Advances in neural information processing systems* 32 (2019).
- [6] Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, and Furu Wei. 2022. Beats: Audio pre-training with acoustic tokenizers. *arXiv preprint arXiv:2212.09058* (2022).
- [7] Bowen Cheng, Anwesa Choudhuri, Ishan Misra, Alexander Kirillov, Rohit Girdhar, and Alexander G. Schwing. 2021. Mask2Former for Video Instance Segmentation. *arXiv:2112.10764* [cs.CV]
- [8] Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. Don't take the easy way out: Ensemble based methods for avoiding known dataset biases. *arXiv preprint arXiv:1909.03683* (2019).
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [10] Shengyi Gao, Zhe Chen, Guo Chen, Wenhai Wang, and Tong Lu. 2023. AVSegFormer: Audio-Visual Segmentation with Transformer. *arXiv:2307.01146* [cs.CV]
- [11] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 776–780.
- [12] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 6904–6913.
- [13] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. 2021. Embodied intelligence via learning and evolution. *Nature communications* 12, 1 (2021), 5721.
- [14] Dawei Hao, Yuxin Mao, Bowen He, Xiaodong Han, Yuchao Dai, and Yiran Zhong. 2023. Improving audio-visual segmentation with bidirectional generation. *arXiv preprint arXiv:2308.08288* (2023).
- [15] 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. 2017. CNN architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 131–135.
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- [17] Di Hu, Yake Wei, Rui Qian, Weiyao Lin, Ruihua Song, and Ji-Rong Wen. 2021. Class-aware sounding objects localization via audiovisual correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 12 (2021), 9844–9859.
- [18] Shaofei Huang, Han Li, Yuqing Wang, Hongji Zhu, Jiao Dai, Jizhong Han, Wenge Rong, and Si Liu. 2023. Discovering Sounding Objects by Audio Queries for Audio Visual Segmentation. *arXiv:2309.09501* [cs.CV]
- [19] Jitesh Jain, Jiachen Li, Mang Tik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. 2023. Oneformer: One transformer to rule universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2989–2998.
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. Segment Anything. *arXiv:2304.02643* (2023).
- [21] Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2023. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. *arXiv preprint arXiv:2311.16922* (2023).
- [22] Kexin Li, Zongxin Yang, Lei Chen, Yi Yang, and Jun Xun. 2023. Cntr: Combinatorial-dependence audio-queried transformer for audio-visual video segmentation. *arXiv preprint arXiv:2309.09709* (2023).
- [23] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355* (2023).
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*. 2980–2988.
- [25] Yuhang Ling, Yuxi Li, Zhenye Gan, Jiangning Zhang, Mingmin Chi, and Yabiao Wang. 2023. Hear to Segment: Unmixing the Audio to Guide the Semantic Segmentation. *arXiv:2305.07223* [cs.SD]
- [26] Chen Liu, Peike Li, Xingqun Qi, Hu Zhang, Lincheng Li, Dadong Wang, and Xin Yu. 2023. Audio-Visual Segmentation by Exploring Cross-Modal Mutual Semantics. *arXiv:2307.16620* [cs.SD]
- [27] Chen Liu, Peike Li, Hu Zhang, Lincheng Li, Zi Huang, Dadong Wang, and Xin Yu. 2023. BAVS: Bootstrapping Audio-Visual Segmentation by Integrating Foundation Knowledge. *arXiv preprint arXiv:2308.10175* (2023).
- [28] Jinxiang Liu, Chen Ju, Chaofan Ma, Yanfeng Wang, Yu Wang, and Ya Zhang. 2023. Audio-aware Query-enhanced Transformer for Audio-Visual Segmentation. *arXiv:2307.13236* [cs.SD]
- [29] Jinxiang Liu, Yu Wang, Chen Ju, Ya Zhang, and Weidi Xie. 2023. Annotation-free Audio-Visual Segmentation. *arXiv preprint arXiv:2305.11019* (2023).
- [30] Shentong Mo and Yapeng Tian. 2023. AV-SAM: Segment anything model meets audio-visual localization and segmentation. *arXiv preprint arXiv:2305.01836* (2023).
- [31] David Nilsson, Aleksis Pirinen, Erik Gärtner, and Cristian Sminchisescu. 2021. Embodied visual active learning for semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 2373–2383.
- [32] Sooyoung Park, Arda Senocak, and Joon Son Chung. 2023. Marginnce: Robust sound localization with a negative margin. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- [33] Sainandan Ramakrishnan, Aishwarya Agrawal, and Stefan Lee. 2018. Overcoming language priors in visual question answering with adversarial regularization. *Advances in Neural Information Processing Systems* 31 (2018).
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28 (2015).
- [35] Ramprasaath R Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Shalini Ghosh, Larry Heck, Dhruv Batra, and Devi Parikh. 2019. Taking a hint: Leveraging explanations to make vision and language models more grounded. In *Proceedings of the IEEE/CVF international conference on computer vision*. 2591–2600.
- [36] Arda Senocak, Tae-Hyun Oh, Junsik Kim, Ming-Hsuan Yang, and In So Kweon. 2018. Learning to localize sound source in visual scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4358–4366.
- [37] Remo MA Van der Heiden, Christian P Janssen, Stella F Donker, Lotte ES Harde- man, Keri Mans, and J Leon Kenemans. 2018. Susceptibility to audio signals during autonomous driving. *PLoS one* 13, 8 (2018), e0201963.
- [38] Yaoting Wang, Weisong Liu, Guangyao Li, Jian Ding, Di Hu, and Xi Li. 2023. Prompting Segmentation with Sound is Generalizable Audio-Visual Source Localizer. *arXiv preprint arXiv:2309.07929* (2023).
- [39] Jiannan Wu, Yi Jiang, Peize Sun, Zehuan Yuan, and Ping Luo. 2022. Language as Queries for Referring Video Object Segmentation. *arXiv:2201.00487* [cs.CV]
- [40] Jialin Wu and Raymond Mooney. 2019. Self-critical reasoning for robust visual question answering. *Advances in Neural Information Processing Systems* 32 (2019).
- [41] Shilin Yan, Renrui Zhang, Ziyu Guo, Wenchao Chen, Wei Zhang, Hongyang Li, Yu Qiao, Zhongjiang He, and Peng Gao. 2023. Referred by Multi-Modality: A Unified Temporal Transformer for Video Object Segmentation. *arXiv:2305.16318* [cs.CV]
- [42] Chao Yang, Su Feng, Dongsheng Li, Huawei Shen, Guoqing Wang, and Bin Jiang. 2021. Learning content and context with language bias for visual question answering. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 1–6.
- [43] Chen Yuanhong, Liu Yuyuan, Wang Hu, Liu Fengbei, Wang Chong, and Carneiro Gustavo. 2023. A Closer Look at Audio-Visual Semantic Segmentation. *arXiv preprint arXiv:2304.02970* (2023).
- [44] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. 2019. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision* 127, 3 (2019), 302–321.
- [45] Jinxing Zhou, Xuyang Shen, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, and Yiran Zhong. 2023. Audio-Visual Segmentation with Semantics. *arXiv preprint arXiv:2301.13190* (2023).
- [46] Jinxing Zhou, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, and Yiran Zhong. 2022. Audio-Visual Segmentation. In *European Conference on Computer Vision*.
- [47] Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. 2024. IB: Alleviating Hallucinations in Large Vision-Language Models via Image-Biased Decoding. *arXiv preprint arXiv:2402.18476* (2024).
- [48] Jannik Zürn and Wolfram Burgard. 2022. Self-Supervised Moving Vehicle Detection from Audio-Visual Cues. *arXiv:2201.12771* [cs.CV]

929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044