

# Probing Cross-modal Information Hubs in Audio-Visual LLMs

Jihoo Jung<sup>1</sup> Chaeyoung Jung<sup>1</sup> Ji-Hoon Kim<sup>2</sup> Joon Son Chung<sup>1</sup>

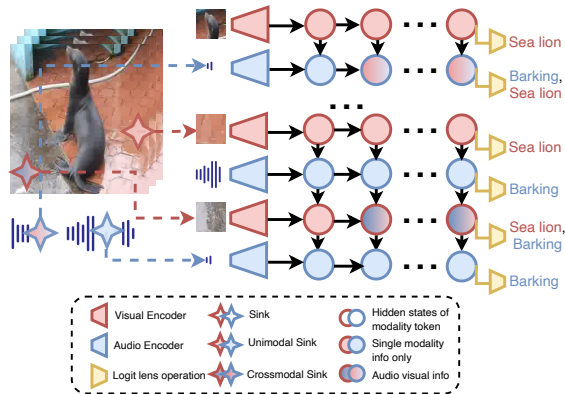
## Abstract

Audio-visual large language models (AVLLMs) have recently emerged as a powerful architecture capable of jointly reasoning over audio, visual, and textual modalities. In AVLLMs, the bidirectional interaction between audio and video modalities introduces intricate processing dynamics, necessitating a deeper understanding of their internal mechanisms. However, unlike extensively studied text-only or large vision language models, the internal workings of AVLLMs remain largely unexplored. In this paper, we focus on cross-modal information flow between audio and visual modalities in AVLLMs, investigating where information derived from one modality is encoded within the token representations of the other modality. Through an analysis of multiple recent AVLLMs, we uncover two common findings. First, AVLLMs primarily encode integrated audio-visual information in sink tokens. Second, sink tokens do not uniformly hold cross-modal information. Instead, a distinct subset of sink tokens, which we term cross-modal sink tokens, specializes in storing such information. Based on these findings, we further propose a simple training-free hallucination mitigation method by encouraging reliance on integrated cross-modal information within cross-modal sink tokens. Our code is available at <https://github.com/kaistmm/crossmodal-hub>.

## 1. Introduction

Recent advancements in large language models (LLMs) have catalyzed the rapid evolution of multimodal LLMs (MLLMs), extending text-centric capabilities to encompass

<sup>1</sup>Department of Electrical Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea <sup>2</sup>The Graduate School of Advanced Imaging Science, Multimedia & Film, Chung-Ang University, Seoul, Republic of Korea. Correspondence to: Joon Son Chung <joonson@kaist.ac.kr>.



**Figure 1. Cross-modal information is primarily stored in cross-modal sink tokens.** Consider an audiovisual clip of a barking sea lion. Cross-modal sink tokens aggregate cues from both modalities, whereas unimodal sink tokens encode information solely from their native modality.

diverse modalities (Yu et al., 2024; Weng et al., 2024; Huang et al., 2024). Among these, audio-visual LLMs (AVLLMs), which integrate auditory and visual inputs via an LLM decoder to generate textual responses, are crucial for achieving a holistic understanding of real-world environments. By jointly processing audio-visual sensory streams, AVLLMs enable reasoning capabilities that closely mirror human multimodal perception, encompassing a comprehensive scope of multimodal context (Zhang et al., 2023; Cheng et al., 2024). As these models continue to evolve and gain widespread adoption, investigating their internal mechanisms has become an imperative step to ensure safety and robustness.

Extensive research has investigated the internal mechanisms of LLMs and LVLMs. Specifically, causal tracing has been utilized to track information transfer (Basu et al., 2024; Kiciman et al., 2024; Meng et al., 2022; Li et al., 2026). In parallel, recent works have employed sparse autoencoders to interpret their internal representations (Huben et al., 2024; Shi et al., 2025; Pach et al., 2025; Lim et al., 2025). However, the internal dynamics of AVLLMs remain largely unexplored. This gap is particularly significant as the integration of the audio modality introduces a unique complexity, necessitating a deeper understanding of simultaneous audio-visual interaction. Elucidating these mechanisms is critical to not only advance model architectures but also verify the factual integrity of models in complex multimodal contexts.

In this work, we explore the internal mechanisms of cross-modal interaction, specifically investigating where informa-

tion derived from one modality (audio or visual) is stored within the token representations of the other modality in AVLLMs. To this end, we adapt the well-established causal tracing technique (Meng et al., 2022; Basu et al., 2024) via our proposed unimodal dominance based framework. By analyzing instances where the model’s output is governed by a single dominant modality, we can pinpoint specific tokens within the non-dominant modality that encapsulate information from the dominant modality.

Our key findings are twofold. First, we demonstrate that attention sink tokens-known to receive disproportionately high attention weights in LLMs and LVLMs-serve as the primary repositories for storing cross-modal information. Second, we find that these sink tokens do not uniformly hold such information. By distinguishing between unimodal sink tokens, which attract substantial attention from their own modality and cross-modal sink tokens, which are heavily attended to by the other modality, we demonstrate that cross-modal sink tokens serve as the distinct carriers of cross-modal information. As illustrated in Fig. 1, cross-modal sink tokens, characterized by blended colors, encapsulate comprehensive audio-visual concepts (e.g., “barking” and “sea lion”), whereas unimodal sink tokens capture information solely from their respective modality.

Finally, leveraging these mechanistic insights, we propose a simple, training-free method to mitigate object hallucinations in AVLLMs. By steering attention toward cross-modal sink tokens, our approach enhances audio-visual integration and significantly reduces hallucinations.

In summary, our contributions are:

- We uncover that cross-modal information in AVLLMs is not uniformly distributed but is localized within attention sink tokens.
- We introduce a functional categorization of sink tokens into unimodal and cross-modal types, revealing that cross-modal sink tokens serve as the critical hubs for integrating audio-visual information.
- We propose a simple, training-free method that mitigates object hallucinations by strategically amplifying the influence of cross-modal sink tokens.

## 2. Related Works

**Audio-visual large language models.** Extending the capabilities of LLMs, AVLLMs have recently emerged, broadening text-centric understanding to encompass audio and visual perception (Zhang et al., 2023; Cheng et al., 2024; Chowdhury et al., 2024; Lyu et al., 2023; Ye et al., 2024; Tang et al., 2025; Guo et al., 2025; Xu et al., 2025a;b). By jointly processing visual and auditory signals, AVLLMs enable context-aware multimodal reasoning

for complex real-world scenarios. Early works such as VideoLLaMA (Zhang et al., 2023) integrated audio-visual inputs into LLaMA (Touvron et al., 2023). Building on this, video-SALMONN (Sun et al., 2024a) enhanced fine-grained temporal processing to improve speech understanding. More recently, Qwen2.5-Omni and Qwen3-Omni (Xu et al., 2025a;b) advanced the field by introducing the capability to generate both text and natural speech as outputs. To effectively fuse audio, visual, and text modalities, various architectural approaches have been explored. Recent state-of-the-art AVLLMs (Xu et al., 2025a;b; Tang et al., 2025; Sun et al., 2025; Guo et al., 2025) process audio and visual inputs via distinct encoders, where the resulting audio and visual embeddings are temporally interleaved and concatenated with text embeddings to serve as input for the LLM backbone. We focus our analysis on this established architecture, aiming to investigate the underlying mechanisms of information exchange between audio and visual modalities.

**Mechanistic interpretability.** Mechanistic interpretability aims to reverse engineer how neural networks process information and make decisions, with the goal of improving transparency, reliability, and trustworthiness. While recent works have leveraged various techniques to probe internal mechanisms, such as causal tracing (Meng et al., 2022; Kiciman et al., 2024; Basu et al., 2024; Li et al., 2026), circuit discovery (Wang et al., 2023; Nanda et al., 2023; He et al., 2025), and sparse autoencoders (Huben et al., 2024; Shi et al., 2025; Pach et al., 2025; Lim et al., 2025) in LLMs and LVLMs, the inner workings of AVLLMs remain largely unexplored. Unlike unimodal or bimodal LLMs, AVLLMs facilitate bidirectional information exchange between audio and visual modalities at the token level. These complex interactions introduce new dynamics, necessitating specialized interpretability analyses.

**Causal tracing.** Among interpretability methods, causal tracing is a widely adopted technique that identifies the causal contribution of specific model components to the final predictions, grounded in causal mediation analysis (Pearl, 2001; Vig et al., 2020). Specifically, the process entails three distinct forward passes: (i) a *clean run* with the original input, (ii) a *corrupted run* with a perturbed input designed to degrade the model’s prediction, and (iii) a *corrupted-with-restoration run*, in which a hidden state from the clean run is patched into the corresponding position of the corrupted run. By quantifying the recovery of the clean prediction, this step determines whether the patched activation encodes critical information, enabling us to trace the causal flow.

## 3. Unimodal Dominance Framework for Tracing Cross-modal Information

To trace bidirectional cross-modal information flow, we first introduce a unimodal dominance framework (Sec. 3.1),

followed by the integration of causal tracing into this setting (Sec. 3.2). Subsequently, we define a metric to quantify the influence of causal states (Sec. 3.3).

### 3.1. Unimodal Dominance Framework

The unimodal dominance framework captures scenarios where a single modality governs the model’s output by providing decisive cues, while its counterpart remains ambiguous. For instance, in identifying a tennis match, the visual modality governs the output by revealing distinct visual features such as racket, whereas the audio modality remains ambiguous as the impact sounds often resemble those of other sports. Positing that information from the dominant modality is propagated into the non-dominant stream via self-attention within transformer block, we aim to localize these propagated signals within the non-dominant stream. Specifically, we leverage audio-dominant samples to identify which visual tokens primarily encode audio-derived information, and video-dominant samples to determine which audio tokens encode video-derived cues.

We classify an instance as *Audio-Dominant* when the joint prediction  $\hat{y}_{av}$  aligns with the audio-only prediction  $\hat{y}_a$  but diverges from the erroneous video-only prediction  $\hat{y}_v$ :

$$\hat{y}_{av} = \hat{y}_a \neq \hat{y}_v. \quad (1)$$

Conversely, the *Video-Dominant* setting is identified when the joint prediction aligns with the video-only prediction but differs from the audio counterpart:

$$\hat{y}_{av} = \hat{y}_v \neq \hat{y}_a. \quad (2)$$

### 3.2. Causal Tracing under Unimodal Dominance Framework

Fig. 2 illustrates causal tracing under unimodal dominance framework with examples. In the *clean run*, the model receives both audio and video inputs without modification, allowing full cross-modal interaction. The model is prompted with a multiple-choice question (e.g., “Which option best describes the clip? A. People farting B. People humming...”) and yields the clean output  $o_{\text{clean}}$ .

To construct the corrupted run, we perturb the dominant modality’s input tokens by zeroing out their raw representations before they are fed into the corresponding modality encoder. As a result, the model is forced to rely solely on the non-dominant modality token set,  $\text{Tokens}_{\text{non-dom}}$ , which encodes misleading semantic information, thereby producing an incorrect prediction  $o_{\text{corrupt}}$ .

We construct a *corrupted-with-restoration run* by patching the hidden states from the clean run ( $h_S^{\text{clean}}$ ) into the corrupted run for a subset of tokens  $S \subset \text{Tokens}_{\text{non-dom}}$ . This patching is applied to the hidden states feeding into

the self-attention sublayer in every Transformer block. Details on the selection of patching locations are provided in Appendix B.1. Patching  $h_S^{\text{clean}}$  is expected to partially recover the clean prediction, as  $h_S^{\text{clean}}$  encodes information originating from the dominant modality via cross-modal interactions.

### 3.3. Metrics for Causal Tracing

We quantify the amount of cross-modal information encoded in the hidden states of the token subset  $S$  by measuring changes in output probabilities between the corrupted run and the restored run. Let  $P[o]$  and  $P_{h_S^{\text{clean}}}[o]$  denote the probability of generating output  $o$  under the corrupted run and the restored run, respectively. Following (Meng et al., 2022), we define the *indirect effect* (IE) with respect to the clean output  $o_{\text{clean}}$  as

$$\text{IE}_{\text{clean}}(S) = P_{h_S^{\text{clean}}}[o_{\text{clean}}] - P[o_{\text{clean}}]. \quad (3)$$

A high  $\text{IE}_{\text{clean}}(S)$  indicates that  $S$  encodes a great amount of information originating from the dominant modality.

We also define the indirect effect with respect to the corrupted output  $o_{\text{corrupt}}$  as

$$\text{IE}_{\text{corrupt}}(S) = P[o_{\text{corrupt}}] - P_{h_S^{\text{clean}}}[o_{\text{corrupt}}]. \quad (4)$$

A high  $\text{IE}_{\text{corrupt}}(S)$  indicates that  $S$  incorporates significant signals from the dominant modality, thereby effectively overriding the ambiguous cues of the non-dominant modality.

Consequently, we utilize these metrics to identify which tokens serve as key mediators of cross-modal information flow; higher values indicate that the subset of token  $S$  acts as a critical repository of cross-modal information.

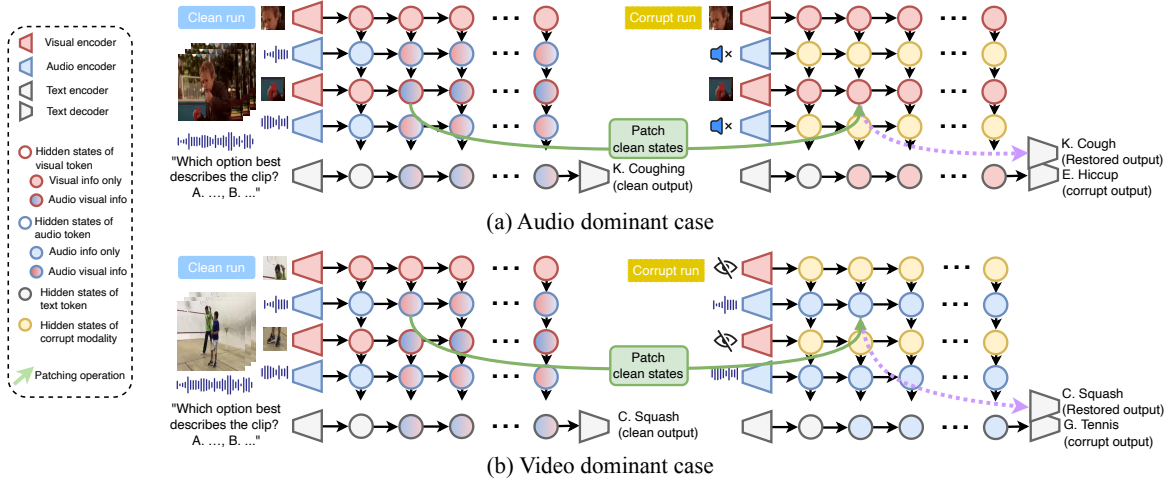
## 4. Where Is Cross-modal Information Located?

Using the causal tracing framework introduced in Sec. 3, we now investigate which tokens primarily encode cross-modal information in AVLLMs.

### 4.1. Hypothesis

Drawing upon recent findings in LVLMs on information storage, we formulate two competing hypotheses: object-centric localization and sink-centric localization

**Object-centric localization.** (Neo et al., 2025) demonstrate that object-centric information is primarily stored in object tokens—tokens corresponding to spatial locations of the object in the original image in LVLMs. This suggests that cross-modal information might be localized within object-aligned tokens, arising from the interaction between the object representation in one modality and the corresponding



**Figure 2. Causal Tracing under the Unimodal Dominance Framework.** For the audio-dominant case, the corrupt run is constructed by corrupting the audio modality, and restoration is conducted by patching hidden states of video tokens from the clean run. Conversely, for the video-dominant case, the video modality is corrupted, and hidden states of audio tokens are patched from the clean run. We expect patching the non-dominant tokens to partially recover the clean prediction, as these tokens encode information transferred from the dominant modality.

object in the complementary modality. Here, we define *audio object tokens* as those capturing the temporal segments of the object’s sound, and *video object tokens* as those encoding its spatio-temporal visual regions.

**Sink-centric localization.** In contrast, (Darcet et al., 2024; Luo et al., 2026) demonstrate that attention sink aggregate high-level global visual information in LVLMs. This suggests that cross-modal information might be localized within sink tokens, acting as the primary locus for abstract semantic cues derived from both modalities. This hypothesis fundamentally contrasts with the object-centric view, as sink tokens are known to emerge from non-object, background positions in LVLMs. For clarity, we refer to sink tokens in the audio modality as *audio sink tokens* and those in the video modality as *video sink tokens*.

## 4.2. Analysis Configuration

To validate these hypotheses, we perform causal tracing experiments across three recent open-source AVLLMs: Qwen2.5-Omni (7B/3B) (Xu et al., 2025a), video-SALMONN-o1 (7B) (Sun et al., 2025), and video-SALMONN2+ (7B/3B) (Tang et al., 2025).

**Dataset.** We use a subset of the VGGSound test set (Chen et al., 2020) to construct audio-dominant and video-dominant evaluation cases. VGGSound is a large-scale audio–visual dataset consisting of in-the-wild videos with class label. Building on findings that certain categories exhibit distinct modality dominance (Jiang et al., 2025), we select two disjoint sets of 20 classes, each with 1,000 videos, as candidate pools for the audio-dominant and video-dominant settings, respectively. For each sample, we present 20 candidate options and ask the model to select the option

that best describes the clip. We then retain only the samples that satisfy Eq. (1) and Eq. (2) for the corresponding dominance setting, and use the filtered set as the final evaluation data for each model. Further details are provided in Appendix A.2.1.

**Token identification.** We employ a recent image segmentation model (Ravi et al., 2025) and a sound event detection model (Wu et al., 2025) to identify video and audio object tokens, respectively. For sink tokens, we adopt the definition from LVLMs (Kang et al., 2025; Luo et al., 2026), where sinks are characterized by abnormally large activations in predefined sink dimensions, with one crucial modification. Prior works determine sink tokens on a per-layer basis and report that sink locations are stable across layers. In contrast, we observe substantial layer-wise variation of sink tokens in AVLLMs (Appendix A.1.1). Consequently, we define global sink tokens as the top- $\frac{|\mathcal{T}|}{N}$  tokens with the highest occurrence frequency across all layers, where  $|\mathcal{T}|$  and  $N$  denote the input sequence length and selection sparsity, respectively.

**Patching strategies.** We measure  $IE_{\text{clean}}$  and  $IE_{\text{corrupt}}$  across four distinct patching scenarios: (1) *All non-dominant modality tokens* ( $\text{Tokens}_{\text{nondom}}$ ), which serves as an upper bound on performance recovery; (2) *Object tokens*, using video object tokens for audio-dominant cases and audio object tokens for video-dominant cases; (3) *Sink tokens*, using video sink tokens for audio-dominant cases and audio sink tokens for video-dominant cases. We report results for  $N \in \{2, 3, 4\}$  to control the number of sink tokens; and (4) *Random tokens*, which serve as a baseline and are sampled to match the number of sink tokens.

### 4.3. Finding1: Cross-modal Information is Primarily Encoded in Sink Tokens

Tab. 1 summarizes the results. Across all models in audio-dominant settings, restoring sink tokens consistently yields larger  $IE_{\text{clean}}$  and  $IE_{\text{corrupt}}$  than restoring object tokens or randomly selected tokens with a comparable number of tokens. Notably, sink tokens maintain substantially higher IE values than random tokens even when the number of sink tokens is small ( $N = 3$  or  $N = 4$ ). In video-dominant settings, although the effect is less pronounced for video-SALMONN2+ due to the small number of patched audio tokens, we observe the same overall trend. These results indicate that, in both directions of information storage, cross-modal information is primarily encoded in sink tokens rather than in object tokens or being uniformly distributed across tokens.

## 5. Are Sink Tokens Homogeneous Cross-modal Information Holders?

Motivated by the finding that visual sink tokens aggregate information from the visual tokens attending to them in LVLMs (Darcey et al., 2024; Luo et al., 2026), we analyze the incoming attention to sink tokens. We specifically investigate the existence of modality-specific attention biases—whether individual sink tokens preferentially aggregate attention from their own modality or the complementary modality. Upon establishing the existence of these biases, we further investigate whether the subset of sink tokens that primarily attract attention from the complementary modality specializes in holding cross-modal information.

### 5.1. Dissecting Sink Tokens

To examine modality-specific attention biases, we first quantify how much incoming attention each sink token receives from video versus audio tokens. Concretely, we define the *Modality Dominance Score* (MDS) for a sink token  $i$  at layer  $l$ , which measures the relative dominance of video attention over audio attention:

$$\text{MDS}_i^l = \frac{\bar{A}_{\text{video},i}^l - \bar{A}_{\text{audio},i}^l}{\bar{A}_{\text{video},i}^l + \bar{A}_{\text{audio},i}^l}. \quad (5)$$

Here,  $\bar{A}_{\text{video},i}^l$  and  $\bar{A}_{\text{audio},i}^l$  denote the mean attention scores received by sink token  $i$  from the video and audio modality tokens at layer  $l$ , respectively.

Fig. 3 visualizes MDS values for audio and video sink tokens for a representative example. This visualization reveals that sink tokens may diverge into two groups: some receive incoming attention primarily from their own modality, while others from the other modality. See Appendix B.1.6 for further analysis on MDS.




Figure 3. **Layer-wise MDS of audio (left) and video (right) sink tokens, sorted by the layer-averaged MDS on Qwen2.5-Omni(7B).** Even within audio sink tokens, some tokens predominantly receive attention from the video modality (audio cross-modal sinks; high MDS), while others receive stronger attention from the audio modality (audio unimodal sinks; low MDS). A similar functional heterogeneity is observed for video sink tokens.

Based on these observations, we investigate whether sink tokens receiving high attention from the complementary modality serve as specialized cross-modal information holders. For a fair comparison, we partition the tokens into two equal-sized groups: unimodal sink tokens, which primarily receive attention from their own modality, and cross-modal sink tokens, which are characterized by high attention from the complementary modality. We then replicate the causal tracing in Sec. 4.2 for each group to quantify their distinct contributions.

### 5.2. Finding2: Cross-modal Sink Tokens Serve as the Primary Carriers of Cross-modal Information

Tab. 2 summarizes the results. Across models and modality dominant settings, restoring cross-modal sink tokens yields substantially high  $IE_{\text{clean}}$  and  $IE_{\text{corrupt}}$  scores, vastly outperforming their unimodal counterparts. Remarkably, results from cross-modal sink patching are comparable to those of the full sink set, implying that the observed causal effect of sink tokens is largely attributable to these cross-modal components. This trend indicates that sink tokens are heterogeneous in their information content: cross-modal sink tokens preferentially store and mediate cross-modal information, whereas unimodal sink tokens do not. Additional analyses are provided in Appendix B.1.

## 6. Application: Mitigating Object Hallucination in AVLLMs

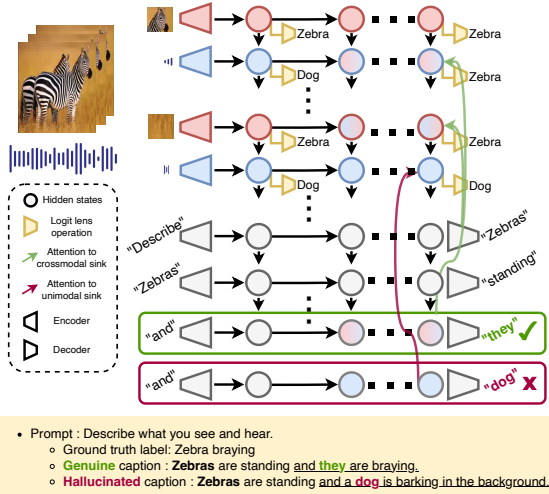
In this section, we analyze object hallucination unique to AVLLMs and propose a mitigation method that leverages our findings on cross-modal sink tokens.

### 6.1. Object Hallucination in AVLLMs

In AVLLMs, unresolved audio–visual disagreement gives rise to a new type of object hallucination. In scenarios where one modality provides a misinterpretation of an object, this conflicting signal may persist across layers, failing to be har-


**Table 1. Patching results across distinct token sets:** all non-dominant modality tokens, object tokens, sink tokens ( $N = 2, 3, 4$ ), and random tokens. We report  $IE_{\text{clean}}$ ,  $IE_{\text{corrupt}}$ , and the number of patched tokens to ensure a fair comparison. **Bold** and underlined values indicate the best and second-best results, respectively; note that the “all non-dominant tokens” case is excluded from this ranking as it serves as an empirical upper bound.

Modality	Ablation	Qwen2.5-Omni(7B)			Qwen2.5-Omni(3B)			video-SALMONN- $\sigma$ 1(7B)			video-SALMONN2+(7B)			video-SALMONN2+(3B)		
		$IE_{\text{clean}} \uparrow$	$IE_{\text{corr}} \uparrow$	#Tokens	$IE_{\text{clean}} \uparrow$	$IE_{\text{corr}} \uparrow$	#Tokens	$IE_{\text{clean}} \uparrow$	$IE_{\text{corr}} \uparrow$	#Tokens	$IE_{\text{clean}} \uparrow$	$IE_{\text{corr}} \uparrow$	#Tokens	$IE_{\text{clean}} \uparrow$	$IE_{\text{corr}} \uparrow$	#Tokens
Audio Dominant	All	9.61	5.28	1440	7.83	3.48	1440	35.55	33.18	1820	6.45	5.27	1210	1.92	2.15	1210
	Object	5.04	2.44	613	3.53	1.12	580	16.22	15.06	852	3.78	3.93	500	0.72	1.16	447
	Sink (N=2)	<b>6.24</b>	<b>2.94</b>	603	<b>6.99</b>	<b>2.70</b>	605	<b>25.33</b>	<b>22.73</b>	818	<b>4.79</b>	<b>4.20</b>	565	<b>1.33</b>	<b>1.38</b>	506
	Sink (N=3)	4.31	1.94	362	6.36	2.08	354	21.42	19.67	514	3.73	3.49	360	0.93	0.94	297
	Sink (N=4)	3.26	1.23	256	5.50	1.64	243	19.10	17.79	364	3.23	3.33	256	0.69	0.65	195
	Random (N=2)	4.24	2.37	603	4.05	1.20	605	20.43	18.11	818	4.21	4.01	565	1.09	0.95	506
	Random (N=3)	2.97	1.55	362	2.71	0.72	354	12.65	13.54	514	3.12	3.51	360	0.67	0.74	297
	Random (N=4)	1.93	0.87	256	1.87	0.65	243	8.70	8.77	364	3.02	3.33	256	0.66	0.53	195
Video Dominant	All	8.21	13.63	249	2.43	8.85	249	3.63	4.08	153	0.46	1.86	60	-0.05	-0.04	60
	Object	<u>4.97</u>	<u>8.44</u>	149	1.59	<u>6.41</u>	149	2.07	0.40	78	<u>0.22</u>	<u>1.71</u>	7	<u>-0.01</u>	-0.06	7
	Sink (N=2)	<b>5.47</b>	<b>8.54</b>	144	<b>2.07</b>	<b>6.87</b>	147	<b>3.57</b>	<b>3.87</b>	117	<b>0.28</b>	<b>2.24</b>	9	-0.02	0.00	29
	Sink (N=3)	4.40	7.12	86	<u>1.62</u>	5.88	109	<u>3.45</u>	<u>3.66</u>	76	0.08	1.70	3	-0.03	-0.06	16
	Sink (N=4)	3.10	6.28	60	1.10	4.78	85	3.30	3.28	52	0.06	1.29	2	<u>-0.01</u>	<b>0.07</b>	13
	Random (N=2)	4.56	6.83	144	1.22	5.29	147	2.86	2.22	117	0.21	<u>1.77</u>	9	-0.02	0.01	29
	Random (N=3)	2.70	3.62	86	0.94	4.34	109	1.69	-0.44	76	0.12	1.28	3	0.00	<u>0.02</u>	16
	Random (N=4)	1.86	2.30	60	0.58	3.46	85	0.9	-2.52	52	0.21	1.28	2	<u>-0.01</u>	<u>0.02</u>	13



**Figure 4. Example of object hallucination in AVLLMs.** While the video modality correctly recognizes the object as a zebra, the audio modality misinterprets the zebra’s braying as a dog’s bark, causing the hallucinated object *dog* to appear in the caption.

monized with the correct modality. Failure to fully suppress the erroneous modality leads to a leakage of incorrect semantic cues, resulting in the model generating captions that incorporate both the correct object and the misinterpreted object. To systematically investigate this phenomenon, we leverage the VGGSound animal category, where such hallucinations frequently occur. Figure 4 illustrates a representative instance: the audio modality misinterprets a zebra’s braying as a dog’s bark, which subsequently manifests as a hallucinated object in the final caption. Consistent with findings by (Nishimura et al., 2024), these hallucinated objects show a tendency to emerge in the second clause, often contextualized as ‘background’ elements.



**Figure 5. Averaged attention mass to cross-modal and unimodal sink tokens across 70 genuine and 70 hallucinated samples.** Genuine object maintain dominant attention on cross-modal sinks across all layers. Conversely, hallucinated object reveal a significant surge in attention to unimodal sinks, occasionally surpassing that of cross-modal sinks.

Delving deeper into these hallucinations, we investigated the model’s attention dynamics during the generation of genuine versus hallucinated objects. To enable a controlled comparison, we restricted our analysis to objects appearing in the second clause. Using the 70 samples obtained for each case, we performed a layer-wise comparison of the attention allocated to cross-modal sink and unimodal sink tokens (See Appendix A.3.1 for the analysis on object tokens). Figure 5 illustrates the proportion of attention mass captured by each sink type within the total audio-visual attention. When generating genuine objects, attention to cross-modal sink tokens remains consistently higher than that to unimodal sink tokens across all layers. In contrast, for hallucinated objects, the attention directed toward unimodal sink tokens becomes more prominent, occasionally surpassing that of cross-modal sink tokens. These results suggest that excessive attention to unimodal sinks—which are devoid of cross-modal cues—can lead to the leakage of misinterpreted information. Since these sinks do not benefit from cross-modal interaction, erroneous signals stored

## Probing Cross-modal Information Hubs in Audio-Visual LLMs

**Table 2. Patching results on unimodal and cross-modal token sets.** We additionally include sink token results for a comprehensive baseline. We report  $\text{IE}_{\text{clean}}$ ,  $\text{IE}_{\text{corrupt}}$ , and the number of patched tokens to ensure a fair comparison. **Bold** and underlined values indicate the best and second-best results, respectively, compared within each block of the same patch size ( $N$ ) across sink, unimodal, and cross-modal tokens.

Modality	Ablation	Qwen2.5-Omni(7B)			Qwen2.5-Omni(3B)			video-SALMONN-o1(7B)			video-SALMONN2+(7B)			video-SALMONN2+(3B)		
		$\text{IE}_{\text{clean}} \uparrow$	$\text{IE}_{\text{corr}} \uparrow$	#Tokens	$\text{IE}_{\text{clean}} \uparrow$	$\text{IE}_{\text{corr}} \uparrow$	#Tokens	$\text{IE}_{\text{clean}} \uparrow$	$\text{IE}_{\text{corr}} \uparrow$	#Tokens	$\text{IE}_{\text{clean}} \uparrow$	$\text{IE}_{\text{corr}} \uparrow$	#Tokens	$\text{IE}_{\text{clean}} \uparrow$	$\text{IE}_{\text{corr}} \uparrow$	#Tokens
Audio Dominant	Sink (N=2)	<b>6.24</b>	<b>2.94</b>	603	<b>6.99</b>	<b>2.70</b>	605	<b>25.33</b>	<b>22.73</b>	818	<b>4.79</b>	<b>4.20</b>	565	<b>1.33</b>	<b>1.38</b>	506
	Unimodal (N=2)	0.65	0.23	301	0.89	0.31	302	7.25	6.82	409	2.32	3.03	282	0.21	0.45	252
	Crossmodal (N=2)	<u>5.58</u>	<u>2.95</u>	301	<u>6.57</u>	<u>2.33</u>	302	<u>21.30</u>	<u>19.93</u>	409	<u>4.16</u>	<u>3.69</u>	282	<u>1.27</u>	<u>1.14</u>	252
	Sink (N=3)	<b>4.31</b>	<b>1.94</b>	362	<b>6.36</b>	<b>2.08</b>	354	<b>21.42</b>	<b>19.67</b>	514	<b>3.73</b>	<b>3.49</b>	360	<b>0.93</b>	<b>0.94</b>	297
	Unimodal (N=3)	0.92	0.39	181	1.02	0.18	177	7.02	7.03	257	2.06	2.87	180	0.19	0.42	148
	Crossmodal (N=3)	<u>3.54</u>	<u>1.52</u>	181	<u>5.73</u>	<u>1.85</u>	177	<u>16.81</u>	<u>15.78</u>	257	<u>3.35</u>	<u>3.20</u>	180	<u>0.77</u>	<u>0.70</u>	148
	Sink (N=4)	<b>3.26</b>	<b>1.23</b>	256	<b>5.50</b>	<b>1.64</b>	243	<b>19.10</b>	<b>17.79</b>	364	<b>3.23</b>	<b>3.33</b>	256	<u>0.69</u>	<b>0.65</b>	195
	Unimodal (N=4)	0.71	0.36	128	1.07	0.32	121	6.19	6.10	182	2.11	2.78	128	0.21	0.36	97
	Crossmodal (N=4)	<u>2.70</u>	<u>0.99</u>	128	<u>4.90</u>	<u>1.28</u>	121	<u>14.24</u>	<u>13.95</u>	182	<u>2.82</u>	<u>3.05</u>	128	<b>0.76</b>	<u>0.64</u>	97
	Sink (N=2)	<b>5.47</b>	<b>8.54</b>	144	<b>2.07</b>	<b>6.87</b>	147	<b>3.57</b>	<b>3.87</b>	117	<b>0.28</b>	<b>2.24</b>	9	<u>-0.02</u>	0.00	29
	Unimodal (N=2)	1.93	3.54	72	0.35	3.43	73	-0.01	-5.00	58	0.18	1.31	4	<b>0.00</b>	<b>0.03</b>	14
	Crossmodal (N=2)	<u>3.03</u>	<u>4.53</u>	72	<u>1.25</u>	<u>4.48</u>	73	<u>3.53</u>	<u>3.72</u>	58	<u>0.26</u>	<u>2.19</u>	4	<u>-0.02</u>	<u>0.01</u>	14
Sink (N=3)	<b>4.40</b>	<b>7.12</b>	86	<b>1.62</b>	<b>5.88</b>	109	<b>3.45</b>	<b>3.66</b>	76	<b>0.08</b>	<b>1.70</b>	3	<u>-0.03</u>	<u>-0.06</u>	16	
Unimodal (N=3)	1.72	3.19	43	0.31	3.15	54	0.13	-4.57	38	0.08	1.44	1	0.00	<b>0.06</b>	8	
Crossmodal (N=3)	<u>2.15</u>	<u>3.70</u>	43	<u>1.01</u>	<u>4.11</u>	54	<u>3.30</u>	<u>3.15</u>	38	<b>0.20</b>	<u>1.60</u>	1	<b>0.01</b>	<b>0.06</b>	8	
Sink (N=4)	<b>3.10</b>	<b>6.28</b>	60	<b>1.10</b>	<b>4.78</b>	85	<b>3.30</b>	<b>3.28</b>	52	<u>0.06</u>	<u>1.29</u>	2	<u>-0.01</u>	<b>0.07</b>	13	
Unimodal (N=4)	1.27	2.80	30	0.24	2.77	42	0.18	-4.46	26	<b>0.08</b>	<b>1.33</b>	1	<u>-0.02</u>	<u>-0.02</u>	6	
Crossmodal (N=4)	<u>1.45</u>	<u>3.02</u>	30	<u>0.63</u>	<u>3.57</u>	42	<u>3.00</u>	<u>2.56</u>	26	0.07	1.25	1	<b>0.02</b>	<u>0.04</u>	6	

within them remain unresolved and subsequently manifest as a hallucinated object in the generated caption.

### 6.2. Adaptive Sink-Guided Decoding

Building on the insight from Section 6.1 that improper attention allocation can lead to object hallucinations, we propose Adaptive Sink-Guided Decoding (ASD) to mitigate this issue. ASD dynamically adjusts the attention weights assigned to cross-modal and unimodal sink tokens during the generation process.

Specifically, we identify global sink tokens and categorize them into cross-modal ( $\mathcal{S}_{\text{cross}}$ ) and unimodal ( $\mathcal{S}_{\text{uni}}$ ) types based on their layer-averaged MDS. Subsequently, at each decoding step  $t$ , we perform two parallel forward passes: (i) the original model forward pass, and (ii) a *calibrated* forward pass where attention to cross-modal and unimodal sink tokens is rebalanced.

In the calibrated pass, we amplify the attention toward cross-modal sink tokens while suppressing that of unimodal sink tokens as follows:

$$\tilde{A}_{t,j} \leftarrow A_{t,j} + \alpha |A_{t,j}|, \quad j \in \mathcal{S}_{\text{cross}}, \quad (6)$$

$$\tilde{A}_{t,j} \leftarrow A_{t,j} - \alpha |A_{t,j}|, \quad j \in \mathcal{S}_{\text{uni}}, \quad (7)$$

where  $A_{t,j}$  denotes the raw attention weights of the token  $j$  at decoding step  $t$ , and  $\alpha$  controls the modulation magnitude, set to 0.6 in all experiments.

We leverage the calibrated pass to steer generation, formulating decoding as an adaptive linear combination in log-probability space between the original and calibrated passes. The  $t$ -th token is thus sampled from the adjusted

distribution:

$$\log \tilde{P}(y_t | \mathbf{x}, y_{<t}) = \gamma_t \log P_{\text{cali}}(y_t | \mathbf{x}, y_{<t}) + (1 - \gamma_t) \log P_{\text{orig}}(y_t | \mathbf{x}, y_{<t}), \quad (8)$$

where  $\mathbf{x}$  represents the multimodal input,  $y_{<t}$  denotes the sequence of preceding tokens,  $P_{\text{cali}}$  and  $P_{\text{orig}}$  denote the probability distributions obtained from the calibrated and original forward passes, respectively, and  $\gamma_t$  serves as the adaptive guidance scale.

Crucially,  $\gamma_t$  is determined adaptively at each step to prioritize the calibrated distribution when the risk of hallucination is high. We identify such ‘‘hallucination-prone’’ steps by detecting excessive attention to unimodal sinks relative to cross-modal sinks in the original forward pass. Accordingly, we compute the base guidance scale  $\gamma_t^{\text{base}}$  as the proportion of total sink attention attributed to unimodal sinks:

$$\gamma_t^{\text{base}} = \frac{\bar{A}_{t,\text{uni}}}{\bar{A}_{t,\text{uni}} + \bar{A}_{t,\text{cross}}}, \quad (9)$$

where  $\bar{A}_{t,\text{uni}}$  and  $\bar{A}_{t,\text{cross}}$  denote the average attention received by unimodal and cross-modal sink tokens, respectively, at decoding step  $t$ . Finally, to ensure stability and prevent abrupt shifts, we apply quadratic soft gating and momentum-based temporal smoothing to  $\gamma_t^{\text{base}}$  to obtain final guidance scale  $\gamma_t$  (see Appendix A.3.2 for details).

### 6.3. Experimental Setting

To assess the efficacy of our approach, we conduct experiments using two representative AVLLMs from Qwen and video-SALMONN families: Qwen2.5-Omni (7B) and video-SALMONN-o1 (7B).

**Datasets.** We evaluate our methods on three datasets. First, to verify whether our approach successfully mitigates the targeted hallucination cases, we use the clean VGGSound-Animal subset, comprising 360 samples, introduced in Sec. 6.1. Additionally, to verify that our method maintains performance on general benchmarks, we evaluate on the standard VGGSound and AudioSet (Gemmeke et al., 2017) datasets. AudioSet is a large-scale audio-visual dataset containing video clips with labeled sound events. For these benchmarks, we utilize approximately 1,000 and 700 clean samples, respectively (see Appendix A.2.2 for details). As the annotations in these datasets are limited to the primary sound source, we augment the ground truth using an object detection model (Carion et al., 2020) to identify additional visible entities. We thus establish the ground truth by combining the labeled source with these detected visible objects.

**Evaluation metrics.** To evaluate object hallucination in the captioning task, we adopt the CHAIR (Rohrbach et al., 2018) metrics, which measure the proportion of mentioned objects that are absent in the ground-truth annotations. CHAIR reports sentence-level ( $C_S$ ) and instance-level ( $C_I$ ) hallucination rates, defined as

$$C_I = \frac{|\{\text{hallucinated objects}\}|}{|\{\text{all mentioned objects}\}|}, C_S = \frac{|\{\text{captions w/ hallucinated objects}\}|}{|\{\text{all captions}\}|},$$

Given that the standard CHAIR is restricted to fixed MS COCO (Lin et al., 2014) objects, we extend its object vocabulary and synonym lists to align with the specific taxonomy of each dataset. Following (Liu et al., 2024), we report the F1 score to evaluate both the richness and accuracy of generated descriptions.

We further adopt ALOHa (Petryk et al., 2024), an open-vocabulary hallucination evaluation metric to overcome the closed-vocabulary limitation of CHAIR. ALOHa utilizes a large language model (GPT-3.5-turbo) to extract groundable objects from generated captions and computes their semantic similarity to the ground-truth objects derived from reference captions.


**Baselines.** We compare our approach with two training-free hallucination mitigation methods adapted to the audio-visual setting (see Appendix A.4 for adaptation details): (i) PAI (Liu et al., 2024), which modulates the attention mechanism to amplify the contribution of multimodal tokens; (ii) VCD (Leng et al., 2024), which contrasts the output logits derived from original inputs with those from distorted counterparts.

## 6.4. Experimental Results

Tab. 3 presents the results. Whereas PAI and VCD offer negligible gains or occasionally exacerbate hallucinations, ASD delivers consistent and substantial reductions in hal-

**Table 3. Quantitative results of ASD.** We evaluate hallucination using ALOHa and CHAIR, and assess caption richness using F1 scores.

Dataset	Method	Qwen2.5-Omni(7B)				video-SALMONN-o1 (7B)			
		ALOHa $\uparrow$	Cs $\downarrow$	Ci $\downarrow$	F1 $\uparrow$	ALOHa $\uparrow$	Cs $\downarrow$	Ci $\downarrow$	F1 $\uparrow$
VGGSound-Animal	Vanilla	40.71	48.21	37.13	55.24	36.21	37.74	32.09	53.68
	PAI	39.52	51.24	38.11	55.11	36.99	35.26	31.18	53.16
	VCD	40.27	51.52	41.28	52.43	36.40	39.39	33.40	53.37
	ASD	<b>42.77</b>	<b>36.91</b>	<b>34.15</b>	52.44	<b>43.29</b>	<b>25.07</b>	<b>25.71</b>	50.89
VGGSound-All	Vanilla	35.02	30.70	<b>20.67</b>	58.69	32.74	30.63	22.39	53.40
	PAI	34.68	32.21	21.52	58.47	32.44	29.29	22.01	53.15
	VCD	34.60	32.63	22.36	57.09	30.28	30.76	24.31	50.02
	ASD	<b>38.89</b>	<b>29.65</b>	21.74	55.81	<b>36.63</b>	<b>21.11</b>	<b>18.42</b>	50.10
AudioSet	Vanilla	38.24	8.92	10.93	69.73	36.81	11.39	14.91	67.27
	PAI	36.94	11.84	13.09	73.22	36.05	10.95	14.54	67.64
	VCD	36.98	12.28	14.88	71.12	32.50	9.34	12.52	67.74
	ASD	<b>38.32</b>	<b>8.54</b>	<b>10.20</b>	72.98	<b>39.64</b>	<b>6.57</b>	<b>9.50</b>	67.29



(a) Qwen2.5-Omni(7B) (b) video-SALMONN-o1(7B)

Figure 6. Parameter sensitivity of  $\alpha$  with CHAIR metrics.

lucinations. Notably, our improvements are maximized on the VGGSound-Animal benchmark—where AVLLM-specific hallucinations induced by audio-visual disagreement are most prevalent—validating our targeted approach. Crucially, ASD maintains this superiority on the more general VGGSound-All and AudioSet datasets, underscoring its robustness.

**Ablation study on  $\alpha$ .** Fig. 6 illustrates the impact of the modulation magnitude  $\alpha$  on the VGGSound-Animal dataset. We observe that ASD is robust to hyperparameter variations, consistently reducing hallucinations across a wide range of values. That said, we note a trade-off: insufficient  $\alpha$  yields limited hallucination mitigation, while excessive  $\alpha$  degrades the richness of the generated captions. Thus, selecting an optimal  $\alpha$  is essential to balance hallucination suppression with the maintenance of details. See Appendix B.2 for additional analyses and Appendix C for qualitative results.

## 7. Conclusion

In this paper, we investigated the internal mechanisms of AVLLMs to understand where cross-modal information is stored in each modality tokens. To enable tracing of bidirectional information flow between audio and video modalities, we adapt causal tracing via a unimodal dominance framework. Using this framework, we showed that cross-modal information is primarily localized in cross-modal sink tokens. Building on these insights, we proposed a simple, training-free object hallucination mitigation method that steers generation toward cross-modal sink tokens. We hope our findings provide foundational insights for future work on interpreting, diagnosing, and improving AVLLMs.

## Acknowledgement

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (RS-2025-02215122, Development and Demonstration of Lightweight AI Model for Smart Homes).

## Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

## References

- Basu, S., Grayson, M., Morrison, C., Nushi, B., Feizi, S., and Massiceti, D. Understanding information storage and transfer in multi-modal large language models. In *NeurIPS*, 2024.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. End-to-end object detection with transformers. In *Proc. ECCV*, 2020.
- Chen, H., Xie, W., Vedaldi, A., and Zisserman, A. Vgsgound: A large-scale audio-visual dataset. In *Proc. ICASSP*, 2020.
- Chen, T., Chakka, C., Akula, A. R., Thomas, X., and Ghadiyaram, D. Some modalities are more equal than others: Decoding and architecting multimodal integration in mllms. *arXiv:2511.22826*, 2025.
- Cheng, Z., Leng, S., Zhang, H., Xin, Y., Li, X., Chen, G., Zhu, Y., Zhang, W., Luo, Z., Zhao, D., et al. VideoLLaMA 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv:2406.07476*, 2024.
- Chowdhury, S., Nag, S., Dasgupta, S., Chen, J., Elhoseiny, M., Gao, R., and Manocha, D. Meerkat: Audio-visual large language model for grounding in space and time. In *Proc. ECCV*, 2024.
- Darcet, T., Oquab, M., Mairal, J., and Bojanowski, P. Vision transformers need registers. In *Proc. ICLR*, 2024.
- Gemmeke, J. F., Ellis, D. P., Freedman, D., Jansen, A., Lawrence, W., Moore, R. C., Plakal, M., and Ritter, M. Audio set: An ontology and human-labeled dataset for audio events. In *Proc. ICASSP*, 2017.
- Guo, Y., Ma, S., Ma, S., Bao, X., Xie, C.-W., Zheng, K., Weng, T., Sun, S., Zheng, Y., and Zou, W. Aligned better, listen better for audio-visual large language models. In *Proc. ICLR*, 2025.
- He, Y., Zheng, W., Dong, Y., Zhu, Y., Chen, C., and Li, J. Towards global-level mechanistic interpretability: A perspective of modular circuits of large language models. In *Proc. ICML*, 2025.
- Huang, R. et al. Audiogpt: Understanding and generating speech, music, sound, and talking head. In *Proc. AAAI*, 2024.
- Huben, R., Cunningham, H., Smith, L. R., Ewart, A., and Sharkey, L. Sparse autoencoders find highly interpretable features in language models. In *Proc. ICLR*, 2024.
- Jiang, X., Wu, J., Choudhari, V., and Mesgarani, N. Bridging ears and eyes: Analyzing audio and visual large language models to humans in visible sound recognition and reducing their sensory gap via cross-modal distillation. *arXiv:2505.06803*, 2025.
- Jung, C., Jang, Y., Choi, J., and Chung, J. S. Fork-merge decoding: Enhancing multimodal understanding in audio-visual large language models. *arXiv preprint arXiv:2505.20873*, 2025.
- Jung, C., Jang, Y., and Chung, J. S. AVCD: Mitigating hallucinations in audio-visual large language models through contrastive decoding. In *NeurIPS*, 2026.
- Kang, S., Kim, J., Kim, J., and Hwang, S. J. See What You Are Told: Visual attention sink in large multimodal models. In *Proc. ICLR*, 2025.
- Kiciman, E., Ness, R., Sharma, A., and Tan, C. Causal reasoning and large language models: Opening a new frontier for causality. *Transactions on Machine Learning Research*, 2024.
- Leng, S., Zhang, H., Chen, G., Li, X., Lu, S., Miao, C., and Bing, L. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proc. CVPR*, 2024.
- Li, Q., Ye, Z., Feng, X., Zhong, W., Ma, W., and Feng, X. Causal tracing of object representations in large vision language models: Mechanistic interpretability and hallucination mitigation. In *Proc. AAAI*, 2026.
- Lim, H., Choi, J., Choo, J., and Schneider, S. Sparse autoencoders reveal selective remapping of visual concepts during adaptation. In *Proc. ICLR*, 2025.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *Proc. ECCV*, 2014.
- Liu, S., Zheng, K., and Chen, W. Paying more attention to image: A training-free method for alleviating hallucination in vlms. In *Proc. ECCV*, 2024.

- Luo, J., Fan, W.-C., Wang, L., He, X., Rahman, T., Abolmaesumi, P., and Sigal, L. To sink or not to sink: Visual information pathways in large vision-language models. In *Proc. ICLR*, 2026.
- Lyu, C., Wu, M., Wang, L., Huang, X., Liu, B., Du, Z., Shi, S., and Tu, Z. Macaw-llm: Multi-modal language modeling with image, audio, video, and text integration. *arXiv:2306.09093*, 2023.
- Meng, K., Bau, D., Andonian, A., and Belinkov, Y. Locating and editing factual associations in gpt. *NeurIPS*, 2022.
- Nanda, N., Chan, L., Lieberum, T., Smith, J., and Steinhardt, J. Progress measures for grokking via mechanistic interpretability. In *Proc. ICLR*, 2023.
- Neo, C., Ong, L., Torr, P., Geva, M., Krueger, D., and Barez, F. Towards interpreting visual information processing in vision-language models. In *Proc. ICLR*, 2025.
- Nishimura, T., Nakada, S., and Kondo, M. On the audio hallucinations in large audio-video language models. *arXiv:2401.09774*, 2024.
- Pach, M., Karthik, S., Bouniot, Q., Belongie, S., and Akata, Z. Sparse autoencoders learn monosemantic features in vision-language models. In *NeurIPS*, 2025.
- Pearl, J. Direct and indirect effects. *Probabilistic and Causal Inference: The Works of Judea Pearl*, pp. 373, 2001.
- Petryk, S., Chan, D. M., Kachinthaya, A., Zou, H., Canny, J., Gonzalez, J. E., and Darrell, T. ALOHa: A new measure for hallucination in captioning models. In *Proc. NAACL*, 2024.
- Ravi, N., Gabeur, V., Hu, Y.-T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al. SAM 2: Segment anything in images and videos. In *Proc. ICLR*, 2025.
- Rohrbach, A., Hendricks, L. A., Burns, K., Darrell, T., and Saenko, K. Object hallucination in image captioning. In *Proc. EMNLP*, 2018.
- Shi, W., Li, S., Liang, T., Wan, M., Ma, G., Wang, X., and He, X. Route sparse autoencoder to interpret large language models. In *Proc. EMNLP*, 2025.
- Su, J., Chen, J., Li, H., Chen, Y., Qing, L., and Zhang, Z. Activation steering decoding: Mitigating hallucination in large vision-language models through bidirectional hidden state intervention. In *Proc. ACL*, 2025.
- Sun, G., Yu, W., Tang, C., Chen, X., Tan, T., Li, W., Lu, L., MA, Z., Wang, Y., and Zhang, C. video-SALMONN: Speech-enhanced audio-visual large language models. In *Proc. ICML*, 2024a.
- Sun, G., Yang, Y., Zhuang, J., Tang, C., Li, Y., Li, W., MA, Z., and Zhang, C. video-SALMONN-o1: Reasoning-enhanced audio-visual large language model. In *Proc. ICML*, 2025.
- Sun, M., Chen, X., Kolter, J. Z., and Liu, Z. Massive activations in large language models. In *Proc. COLM*, 2024b.
- Tang, C., Li, Y., Yang, Y., Zhuang, J., Sun, G., Li, W., Ma, Z., and Zhang, C. video-SALMONN 2: Captioning-enhanced audio-visual large language models. *arXiv:2506.15220*, 2025.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv:2302.13971*, 2023.
- Vig, J., Gehrmann, S., Belinkov, Y., Qian, S., Nevo, D., Sakenis, S., Huang, J., Singer, Y., and Shieber, S. Causal mediation analysis for interpreting neural nlp: The case of gender bias. *arXiv:2004.12265*, 2020.
- Wang, K., Variengien, A., Conmy, A., Shlegeris, B., and Steinhardt, J. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. In *Proc. ICLR*, 2023.
- Weng, Y., Han, M., He, H., Chang, X., and Zhuang, B. Longvlm: Efficient long video understanding via large language models. In *Proc. ECCV*, 2024.
- Wu, Y., Tsirigotis, C., Chen, K., Huang, C.-Z. A., Courville, A., Nieto, O., Seetharaman, P., and Salamon, J. FLAM: Frame-wise language-audio modeling. In *Proc. ICML*, 2025.
- Xu, J., Guo, Z., He, J., Hu, H., He, T., Bai, S., Chen, K., Wang, J., Fan, Y., Dang, K., et al. Qwen2. 5-omni technical report. *arXiv:2503.20215*, 2025a.
- Xu, J., Guo, Z., Hu, H., Chu, Y., Wang, X., He, J., Wang, Y., Shi, X., He, T., Zhu, X., Lv, Y., Wang, Y., Guo, D., Wang, H., Ma, L., Zhang, P., Zhang, X., Hao, H., Guo, Z., Yang, B., Zhang, B., Ma, Z., Wei, X., Bai, S., Chen, K., Liu, X., Wang, P., Yang, M., Liu, D., Ren, X., Zheng, B., Men, R., Zhou, F., Yu, B., Yang, J., Yu, L., Zhou, J., and Lin, J. Qwen3-omni technical report. *arXiv:2509.17765*, 2025b.
- Ye, Q., Yu, Z., Shao, R., Xie, X., Torr, P., and Cao, X. Cat: Enhancing multimodal large language model to answer questions in dynamic audio-visual scenarios. In *Proc. ECCV*, 2024.
- Yu, T. et al. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *Proc. CVPR*, 2024.

Zhang, H., Li, X., and Bing, L. Video-llama: An instruction-tuned audio-visual language model for video understanding. In *Proc. EMNLP*, 2023.

Zverev, D., Wiedemer, T., Prabhu, A., Bethge, M., Brendel, W., and Koepke, A. Vggsounder: Audio-visual evaluations for foundation models. In *Proc. ICCV*, 2025.

## Appendix Outline

---

The appendix provides detailed implementations, additional analyses, and qualitative examples supporting the main paper. The structure is organized as follows:

### A. Implementation Details

- **A.1. Sink Tokens**
  - A.1.1. Definition of Sink Tokens
  - A.1.2. Selection of Sink Dimension
- **A.2. Dataset**
  - A.2.1. Datasets for Causal Tracing
  - A.2.2. Datasets for Object Hallucination Mitigation Evaluation
- **A.3. Adaptive Sink-Guided Decoding Algorithm**
  - A.3.1. Insignificance of Object Tokens in Hallucination Detection
  - A.3.2. Details of Stability Functions
- **A.4. Details of Adapting Baseline Approaches to Audio-Visual Setting**

### B. Additional Analysis

- **B.1. Causal Tracing Analysis**
  - B.1.1. Bottom-up Analysis
  - B.1.2. Patching Locations Analysis
  - B.1.3. Causal Tracing Results with Different Patching Locations
  - B.1.4. Layerwise Analysis
  - B.1.5. Alternative Corruption Methods Analysis
  - B.1.6. MDS Statistics
- **B.2. Object Hallucination Reduction**
  - B.2.1. Object Hallucination Case Analyses
  - B.2.2. Additional Results on Different Models
  - B.2.3. Reverse ASD as a Counterfactual Test
  - B.2.4. Comparison with Inference-Time Interventions in AVLLMs
  - B.2.5. Latency Overhead of ASD

### C. Qualitative Results

### D. Limitations

---

## A. Implementation Details

### A.1. Sink Tokens

#### A.1.1. DEFINITION OF SINK TOKENS

Sink tokens are tokens that receive disproportionately high attention, exhibit abnormally large activations in specific hidden dimensions (so-called *sink dimensions*  $\mathcal{D}_{\text{sink}}$ ) (Kang et al., 2025), and have high feature norms. Prior studies have characterized sink tokens using diverse criteria, including attention aggregation patterns and activation-based signatures (Darcet et al., 2024; Kang et al., 2025; Sun et al., 2024b).

Following (Kang et al., 2025), we identify sink tokens at each layer based on their activations in predefined sink dimensions. Formally, for each layer  $l$ , the set of sink tokens is defined as

$$\hat{\mathcal{I}}^l = \{j \in \mathcal{I} \mid \phi(\mathbf{x}_j^{l-1}) \geq \tau\}, \quad (10)$$

where  $\mathbf{x}_j^{l-1}$  denotes the pre-attention hidden representation of token  $j$  at layer  $l$ ,  $\tau$  is a predefined threshold, and  $\phi(\cdot)$  is a sink characteristic function defined as

$$\phi(\mathbf{x}_j^{l-1}) = \max_{d \in \mathcal{D}_{\text{sink}}} |\text{RMSNorm}(\mathbf{x}_j^{l-1})[d]|. \quad (11)$$

Here,  $\text{RMSNorm}(\cdot)$  denotes root-mean-square normalization, and  $\mathcal{D}_{\text{sink}}$  is the set of sink dimensions determined from the backbone LLM.

(Kang et al., 2025) report that visual sink tokens in VLMs emerge in early layers and persist across deeper layers. In contrast, we observe substantial variability in sink token identities across layers in audio-visual LLMs. Figure 7 shows the layer-wise average number of sink tokens in Qwen2.5-Omni (7B). The number of sink tokens varies significantly across layers and increases sharply in deeper layers, where a large fraction of tokens become classified as sink tokens. This suggests that sink behavior in AVLLMs is highly dynamic and layer-dependent, likely reflecting hierarchical multimodal information aggregation.

Motivated by this observation, we introduce a *global sink token* definition that accounts for cross-layer variability. Instead of using layer-wise sink sets  $\hat{\mathcal{I}}^l$ , we measure how frequently each token is classified as a sink token across layers. Specifically, for each token  $j$ , we define its sink frequency score as

$$s_j = \sum_{l=1}^L \mathbb{I}[j \in \hat{\mathcal{I}}^l], \quad (12)$$

where  $L$  is the total number of layers and  $\mathbb{I}(\cdot)$  is the indicator function.

We then sort tokens by  $s_j$  and define the top- $K$  tokens as *global sink tokens*, where  $K$  is set to a fixed fraction of the maximum input length:

$$\mathcal{I}_{\text{global}} = \text{TopK} \left( \{s_j\}_{j \in \mathcal{I}}, K = \frac{|\mathcal{T}|}{N} \right), \quad (13)$$

with  $|\mathcal{T}|$  denoting the input token length and  $N$  a normalization constant. This definition yields a stable set of sink tokens that persist across layers and avoids over-selecting transient layer-specific sinks.

#### A.1.2. SELECTION OF SINK DIMENSION

(Kang et al., 2025) reported that, even after multimodal fine-tuning, the dimensions exhibiting massive activations remain largely consistent with those of the base LLM. Specifically, visual sink tokens in VLMs exhibit massive activation along the same dimensions as the BOS token in the base LLM.

Following this observation, we select sink dimensions based on BOS-token activations in the backbone LLM. For Qwen2.5-Omni (7B/3B) and video-SALMONN2+ (7B/3B), whose backbones are Qwen2.5-VL (7B/3B), we identify the BOS token’s massive-activation dimensions in the corresponding base model and use these dimensions as  $\mathcal{D}_{\text{sink}}$  for all subsequent analyses.


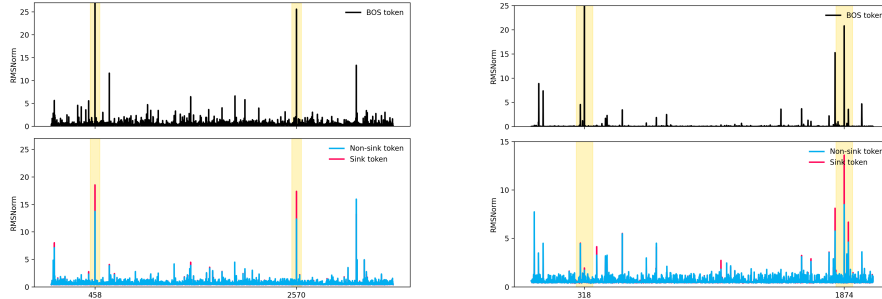


Figure 7. Layer-wise average number of sink tokens across 100 samples in Qwen2.5-Omni (7B). The shaded region denotes one standard deviation. The number of sink tokens vary substantially across layers.

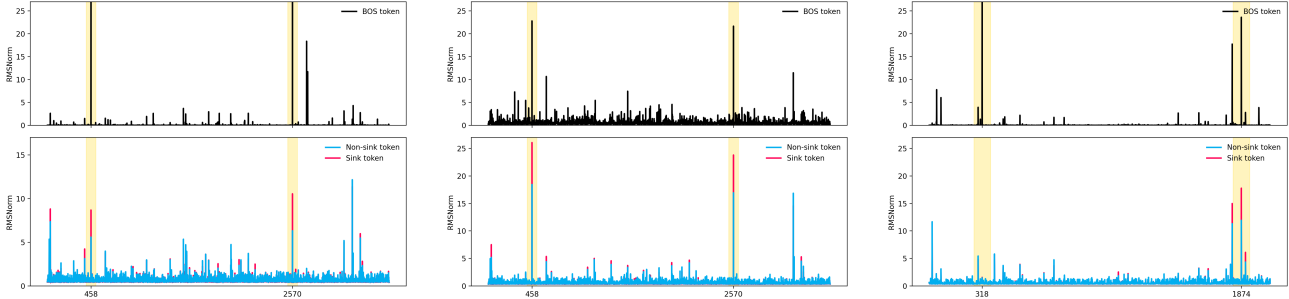
Table 4. Model-specific sink dimensions  $\mathcal{D}_{\text{sink}}$  and thresholds  $\tau$  used to identify layer-wise sink tokens.

Model	$\mathcal{D}_{\text{sink}}$	$\tau$
Qwen2.5-Omni (7B)	(458, 2570)	25
Qwen2.5-Omni (3B)	(318, 1874)	20
video-SALMONN-o1 (7B)	(458, 2570)	25
video-SALMONN2+ (7B)	(458, 2570)	25
video-SALMONN2+ (3B)	(318, 1874)	20



(a) Qwen2.5-Omni (7B)

(b) Qwen2.5-Omni (3B)



(c) video-SALMONN-o1 (7B)

(d) video-SALMONN2+ (7B)

(e) video-SALMONN2+ (3B)

Figure 8. Dimension-wise RMSNorm magnitudes for BOS, sink, and non-sink tokens. Top: Qwen-based models. Bottom: SALMONN-based models.

For video-SALMONN-o1, which is based on Qwen2-7B, we directly adopt the sink-dimension set  $\mathcal{D}_{\text{sink}} = \{458, 2570\}$  reported in (Kang et al., 2025). Table 4 summarizes the sink dimensions and thresholds used in this work.

Figure 8 compares the dimension-wise RMSNorm magnitudes for BOS, sink, and non-sink tokens, averaged over 100 samples. Sink tokens consistently exhibit larger activations along the identified sink dimensions, validating our sink-dimension selection strategy across models.

## A.2. Dataset

### A.2.1. DATASETS FOR CAUSAL TRACING

To analyze the localization of cross-modal information via causal tracing, we construct evaluation subsets based on the *Audio Dominant* and *Video Dominant* cases defined in Sec. 3.1. Following the observations of (Jiang et al., 2025), audio-dominant cases predominantly correspond to human actions with distinctive acoustic cues, whereas video-dominant cases cover sports and visually salient activities.

Based on these findings, we curate two disjoint sets of categories from the VGGSound dataset. For each case, we first collect representative categories explicitly mentioned in prior work and then expand them with semantically aligned categories from VGGSound. To avoid ambiguity in dominance attribution, we exclude overlapping or semantically redundant categories. For example, categories such as *people eating* and *eating with cutlery* are treated as duplicates, and only one is retained.

**Audio dominant categories.** The final set of audio-dominant categories consists of the following 20 classes: people farting, people finger snapping, people slapping, people battle cry, civil defense siren, people sobbing, people whispering, people humming, people shuffling, police radio chatter, people burping, people nose blowing, smoke detector beeping, people booing, people whistling, people eating, people gargling, people coughing, people hiccup, and francolin calling.

**Video dominant categories.** The final set of video-dominant categories consists of the following 20 classes: playing volleyball, playing darts, playing table tennis, skiing, striking bowling, skateboarding, rope skipping, golf driving, playing tennis, mouse clicking, basketball bounce, swimming, slot machine, sailing, playing squash, playing hockey, shooting football, playing badminton, bouncing on trampoline, and parrot talking.

Each category in the VGGSound test set contains 50 samples, resulting in 1,000 samples for the audio-dominant set and 1,000 samples for the video-dominant set before filtering. As discussed earlier, we further filter these samples based on model-specific prediction patterns. Specifically, we retain only samples satisfying  $\hat{y}_{av} = \hat{y}_a \neq \hat{y}_v$  for the audio-dominant case and  $\hat{y}_{av} = \hat{y}_v \neq \hat{y}_a$  for the video-dominant case. These filtered subsets are used for causal tracing analysis for each model.

Table 5. Number of samples retained for causal tracing analysis after applying audio-dominant and video-dominant filtering criteria for each model.

Dominance	Qwen2.5-Omni (7B)	Qwen2.5-Omni (3B)	video-SALMONN-o1 (7B)	video-SALMONN2+ (7B)	video-SALMONN2+ (3B)
Audio Dominant	349	315	184	369	141
Video Dominant	548	646	828	808	890

Tab. 5 summarizes the number of samples retained for each model after applying the dominance-based filtering.

#### A.2.2. DATASETS FOR OBJECT HALLUCINATION MITIGATION EVALUATION

Unlike image captioning, where dedicated benchmarks for measuring object hallucination exist (e.g., MS-COCO (Lin et al., 2014)), there is no established benchmark for object hallucination evaluation in audio-visual captioning. To construct a comparable evaluation framework, we leverage audio-visual classification datasets, VGGSound and AudioSet, where audio labels are reliably annotated but visual background objects are not exhaustively labeled. To obtain comprehensive ground-truth object sets, we additionally apply object detection models to identify unlabeled visual objects.

We evaluate our method on three datasets. Our empirical analysis reveals that object hallucinations occur most frequently in animal-related scenarios, where audio-visual correlations are strong but visually ambiguous. Therefore, we first evaluate on a curated VGGSound-Animal benchmark to quantify hallucination mitigation in challenging cases. To assess generalization and ensure that our method does not degrade performance on broader distributions, we additionally evaluate on VGGSound-All and AudioSet.

**VGGSound-Animal.** We construct a curated VGGSound-Animal benchmark based on VGGSounder (Zverev et al., 2025), a comprehensively re-annotated multi-label evaluation set derived from the original VGGSound dataset. We focus on animal-related classes and exclude bird categories due to severe label ambiguity and semantic overlap. We further filter samples by removing instances where the VGGSounder annotations and the original VGGSound labels are inconsistent. From the remaining samples, we select clips in which the target animal class is present in both audio and visual modalities and is the only annotated class in the clip. This results in approximately 360 clean single-label animal samples.

**VGGSound-All.** To evaluate generalization beyond single-label animal scenarios, we construct a VGGSound-All benchmark that includes multi-label samples. We retain only samples where the VGGSounder annotations are consistent with the original VGGSound labels. For each object class, we randomly sample up to 10 clips to balance the class distribution, resulting in approximately 1,200 samples. Multi-label annotations from VGGSounder are treated as ground-truth object sets.

**AudioSet.** Given that AudioSet contains substantial annotation noise and weak cross-modal alignment, we adopt the subset curated by (Chen et al., 2025). This data was obtained through a rigorous filtering process involving ontology pruning, automated cross-modal verification, and human inspection. In total, we employ approximately 680 clean samples.

**Ground-truth object sets.** Following (Petryk et al., 2024), all video clips are truncated to a maximum duration of 10 seconds. We uniformly sample 10 frames per video and perform object detection using DETR (Carion et al., 2020). Detected

objects are used to construct the final ground-truth object sets for hallucination evaluation.

### A.3. Adaptive Sink-Guided Decoding Algorithm

#### A.3.1. INSIGNIFICANCE OF OBJECT TOKENS IN HALLUCINATION DETECTION

While Sec. 6.1 demonstrated the sensitivity of sink tokens to hallucination, we investigate if object tokens share this characteristic. Contrary to sink tokens, Fig. 9 reveals that attention on object tokens shows no significant divergence between genuine and hallucinated cases. This invariance sharply contrasts with the distinct attention disparity observed in sink tokens (Fig. 5).




Figure 9. Comparison of attention weights on object tokens during the generation of genuine versus hallucinated objects. Unlike the unimodal sink, object tokens do not show a significant disparity between the two scenarios.

#### A.3.2. DETAILS OF STABILITY FUNCTIONS

To ensure stability and prevent abrupt shifts caused by fluctuations in the raw coefficient, we apply threshold-based gating and momentum-based temporal smoothing to  $\gamma_t^{\text{base}}$  to obtain the stabilized coefficient  $\gamma_t$ .

To compute  $\gamma_t$ , we first define the base adaptive weight  $\gamma_t^{\text{base}}$ , which quantifies the relative dominance of unimodal versus cross-modal sink attention at decoding step  $t$ :

$$\gamma_t^{\text{base}} = \frac{\bar{A}_{t,\text{uni}}}{\bar{A}_{t,\text{uni}} + \bar{A}_{t,\text{cross}} + \epsilon}, \quad (14)$$

where  $\bar{A}_{t,\text{uni}}$  and  $\bar{A}_{t,\text{cross}}$  denote the average attention mass assigned to unimodal and cross-modal sink tokens, respectively, and  $\epsilon$  is a small constant for numerical stability.

Concretely, the target coefficient is defined as:

$$\hat{\gamma}_t = \begin{cases} 0, & \gamma_t^{\text{base}} < \tau \text{ or } r_t > \rho, \\ \gamma_{\text{max}} \gamma_t^{\text{base}}, & \text{otherwise,} \end{cases} \quad (15)$$

where  $\gamma_{\text{max}}$  is set to 0.6,  $\tau$  is the adaptive gating threshold set to 0.6,  $r_t$  denotes the average attention mass assigned to text tokens, and  $\rho$  is the text-mass threshold set to 0.5.

Finally, we apply momentum-based temporal smoothing:

$$\gamma_t = \beta \gamma_{t-1} + (1 - \beta) \hat{\gamma}_t, \quad (16)$$

where  $\beta$  is the momentum coefficient set to 0.7, which prevents abrupt temporal fluctuations during decoding.

### A.4. Details of Adapting Baseline Approaches to Audio-Visual Setting

In this section, we describe how we adapted existing image-based baselines for the audio-visual setting.

PAI PAI (Liu et al., 2024) aims to mitigate hallucinations by amplifying the attention scores of image tokens globally. The original mechanism is formulated as:

$$\tilde{A}_{t,j} \leftarrow A_{t,j} + \alpha |A_{t,j}|, \quad j \in \mathcal{S}_{\text{image}}, \quad (17)$$

where  $\mathcal{S}_{\text{image}}$  represents the set of image indices. To adapt PAI for our audio-visual setting, we modify the method to increase the attention weights for both audio and video tokens, rather than image tokens alone.

VCD VCD (Leng et al., 2024) reduces statistical priors and hallucinations by contrasting the original output logits with logits derived from distorted visual inputs. While the original VCD applies noise solely to the image modality, we extend this approach to the audio-visual domain. Specifically, we apply noise to both audio and video inputs to generate the distorted logits for contrastive decoding.

## B. Additional Analysis

### B.1. Causal Tracing Analysis

#### B.1.1. BOTTOM-UP ANALYSIS

To further verify whether our pre-selected token categories indeed coincide with the tokens that are causally most relevant to cross-modal information flow, we additionally conduct a bottom-up, token-level causal analysis on 100 samples using Qwen2.5-Omni(7B). Specifically, we repeat the experiments in Sec. 4.2 by applying causal patching to individual tokens rather than predefined groups. After ranking the tokens by their causal effect  $\Delta$ , we measure the proportion of sink and object tokens ( $\mathcal{S} \cup \mathcal{O}$ ) versus the remaining (*Neither*) tokens among the top 5%, 10%, and 20% most influential tokens. As shown in Tab. 6, the majority of the most influential tokens fall into  $\mathcal{S} \cup \mathcal{O}$ . This supports our premise that sink and object tokens are highly meaningful, providing a reasonable basis for our initial hypothesis.

Table 6. Proportion of sink/object tokens ( $\mathcal{S} \cup \mathcal{O}$ ) versus the remaining (*Neither*) tokens among the top- $k\%$  most influential tokens, ranked by their causal effect  $\Delta$ .  $N$  denotes the parameter for defining global sink tokens (Sec. 4.2).

$N$	Group	Video Dom.			Audio Dom.		
		5%	10%	20%	5%	10%	20%
2	$\mathcal{S} \cup \mathcal{O}$	0.88	0.90	0.84	0.71	0.73	0.77
	<i>Neither</i>	0.12	0.10	0.16	0.29	0.27	0.23
3	$\mathcal{S} \cup \mathcal{O}$	0.82	0.84	0.84	0.64	0.66	0.70
	<i>Neither</i>	0.18	0.16	0.16	0.36	0.34	0.30

To further quantify how much cross-modal information resides outside  $\mathcal{S} \cup \mathcal{O}$ , we measure the proportion (%) of the causal effect  $\Delta$  captured by patching the  $\mathcal{S} \cup \mathcal{O}$  group versus all remaining tokens (*Neither*), relative to the effect of patching all tokens. Our results in Tab. 7 indicate that while some cross-modal information may be distributed outside these key tokens, this fraction is minimal. We clarify that we do not claim cross-modal information is *exclusively* stored within our identified hubs; rather, they serve as its disproportionately dominant hubs.

Table 7. Proportion of the causal effect  $\Delta$  captured by the  $\mathcal{S} \cup \mathcal{O}$  group versus the *Neither* group, relative to patching all tokens.

$N$	Group	Video Dom.	Audio Dom.
2	$\mathcal{S} \cup \mathcal{O}$	0.92	0.83
	<i>Neither</i>	0.11	0.20
3	$\mathcal{S} \cup \mathcal{O}$	0.87	0.73
	<i>Neither</i>	0.16	0.29

#### B.1.2. PATCHING LOCATIONS ANALYSIS

Identifying the optimal layer for hidden state patching is a critical component of the causal tracing framework. Prior studies investigating factual association in text-only LLMs and Large Vision-Language Models (LVLMs) typically compare the efficacy of patching immediately after the Self-Attention (SA) mechanism versus after the MLP layers (Meng et al., 2022). These works generally conclude that factual knowledge is predominantly stored within the MLP layers.

However, our research objective differs significantly. Rather than locating static factual storage, we aim to trace the dynamic flow of cross-modal information between audio and video inputs. Given that cross-modal fusion relies on the token-mixing

capabilities of the attention mechanism, we posit that the Self-Attention layers are the primary locus of this interaction.

Furthermore, we critically assess whether to patch *before* or *after* the Self-Attention (SA) operation. While the output states *after* the SA layer represent the result of information fusion, the critical limitation is that the actual transfer of this fused information to the text modality occurs *during* the attention computation. Consequently, if we patch the output of the SA layer, the fused cross-modal information is not effectively attended to by the text tokens within that layer, resulting in a failure to reflect this information in the model’s final prediction.

Therefore, we hypothesize that patching the hidden states *before* the SA layer (i.e., the input to the SA block) is methodologically superior. Although the pre-attention states may theoretically contain less fully integrated cross-modal information compared to the post-attention outputs, patching at this stage is crucial as it allows the attention mechanism to actively propagate the restored information to the text sequence, thereby ensuring it is incorporated into the model’s output.

To validate this hypothesis, we conducted a control experiment distinct from our main analysis. Unlike the main experiments, which patch the hidden states of the non-dominant modality (where the dominant information is only partially integrated, preventing the prediction probability from being fully restored to the clean run level), here we patched the *dominant* modality. In this setting, since the dominant modality dictates the correct answer, successfully patching its information should result in a near-complete recovery of the correct prediction (i.e., a high  $IE_{\text{clean}}$  and  $IE_{\text{corrupt}}$ ).

We compared patching performance at three specific locations: (1) Before Self-Attention, (2) After Self-Attention, and (3) After MLP. We evaluated the restoration of correct answers using the Qwen2.5-Omni (7B) model in Audio-Dominant scenarios. The results, presented in Tab. 8, demonstrate that patching **Before Self-Attention** yields a significantly higher recovery of the correct prediction compared to other locations. Consequently, we decided to patch the hidden states before the self-attention layer for all main experiments.

Table 8. Comparison of Indirect Effects (IE) across different patching locations on Qwen2.5-Omni (7B) under Audio-Dominant settings. Patching all dominant tokens before the Self-Attention layer results in the most effective information recovery.

Patching Location	$IE_{\text{clean}}$	$IE_{\text{corr}}$
Before Self-Attention	63.72	35.22
After Self-Attention	3.01	2.38
After MLP Layer	20.96	11.67

### B.1.3. CAUSAL TRACING RESULTS WITH DIFFERENT PATCHING LOCATIONS

As discussed earlier, although we primarily patch hidden states at the self-attention module, which we find to be the most suitable patching location, we further examine the robustness of our findings with respect to the patching location. Specifically, following prior work (Meng et al., 2022; Basu et al., 2024; Li et al., 2026), we consider two alternative patching points: (i) immediately after the self-attention block, and (ii) immediately after the MLP block.

The causal tracing results obtained by patching hidden states after the self-attention block are reported in Tab. 9, while those obtained by patching hidden states after the MLP block are shown in Tab. 10. Consistent with the main results, we observe that cross-modal information is predominantly concentrated in sink tokens, with cross-modal sink tokens exhibiting the strongest effects.

### B.1.4. LAYERWISE ANALYSIS

Fig. 10 shows layer-wise causal patching results for the audio-dominant setting on Qwen2.5-Omni (7B), where we patch all tokens from the non-dominant modality. We use a sliding window of size 10, patching tokens across 10 consecutive layers at a time. The results indicate that cross-modal information exchange is most pronounced in the middle layers.

### B.1.5. ALTERNATIVE CORRUPTIONS METHOD ANALYSIS

To verify that our results are robust to corruption choices beyond zeroing out, we repeat the experiments in Sec. 4.2 on Qwen2.5-Omni(7B) using two alternative corruption methods: (1) injecting Gaussian noise into the raw input, and (2) replacing the multimodal encoder output embeddings with their mean embedding, following (Neo et al., 2025). As shown in Tab. 11, both alternatives align with our findings.

Probing Cross-modal Information Hubs in Audio-Visual LLMs

Table 9. Causal tracing results obtained by patching hidden states after the self-attention block.

Modality	Ablation	Qwen2.5-Omni (7B)			Qwen2.5-Omni (3B)			video-SALMONN-o1 (7B)			video-SALMONN2+ (7B)			video-SALMONN2+ (3B)		
		IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens
Audio Dominant	All	9.61	5.29	1440	7.83	3.48	1440	35.55	33.18	1820	6.45	5.27	1210	1.92	2.15	1210
	Object	5.05	2.44	613	3.53	1.12	580	16.22	15.06	852	3.78	3.93	500	0.72	1.16	448
	Random (N=2)	4.01	2.34	603	4.02	1.19	605	19.19	17.98	818	4.27	4.09	566	0.94	1.12	506
	Random (N=3)	2.76	1.14	363	2.21	0.50	355	13.01	12.51	514	3.17	3.53	360	0.83	0.77	298
	Random (N=4)	1.93	1.02	256	1.87	0.53	244	9.32	9.20	364	3.05	3.33	257	0.68	0.57	196
	Sink (N=2)	6.24	2.94	603	7.26	2.43	605	25.33	22.73	818	4.79	4.2	566	1.33	1.38	506
	Sink (N=3)	4.31	1.95	363	6.34	2.06	354	21.51	19.67	514	3.73	3.49	360	0.93	0.94	298
	Sink (N=4)	3.27	1.24	256	5.50	1.64	244	19.10	17.79	364	3.23	3.33	257	0.69	0.65	196
	Unimodal (N=2)	0.65	0.24	301	0.89	0.3	302	7.25	6.82	409	2.32	3.03	283	0.21	0.45	253
	Unimodal (N=3)	0.92	0.4	181	1.02	0.18	177	7.02	7.03	257	2.06	2.87	180	0.19	0.42	149
	Unimodal (N=4)	0.71	0.37	128	1.07	0.32	122	6.19	6.10	182	2.11	2.78	128	0.21	0.36	98
	Crossmodal (N=2)	5.59	2.96	302	6.55	2.32	303	21.30	19.93	409	4.16	3.69	283	1.27	1.14	253
	Crossmodal (N=3)	3.55	1.53	182	5.73	1.85	178	16.81	15.78	257	3.35	3.2	180	0.77	0.70	149
	Crossmodal (N=4)	2.7	1.00	128	4.9	1.28	122	14.24	13.95	182	2.82	3.05	129	0.76	0.64	98
	Video Dominant	All	8.21	13.64	249	2.43	8.85	250	3.63	4.09	153	0.46	1.86	60	-0.05	-0.04
Object		4.97	8.44	149	1.59	6.41	149	2.07	0.40	78	0.22	1.71	7	-0.01	-0.06	7
Random (N=2)		4.15	6.61	144	1.23	5.43	148	2.28	1.66	117	0.24	1.81	9	-0.02	0.02	29
Random (N=3)		2.71	4.28	87	0.80	4.20	109	1.67	-0.72	76	0.11	1.4	4	0	0.05	17
Random (N=4)		1.84	3.09	61	0.44	3.25	85	1.03	-2.73	52	0.22	1.28	2	-0.01	-0.01	13
Sink (N=2)		5.47	8.54	144	2.07	6.87	148	3.57	3.87	117	0.28	2.24	9	-0.02	0	29
Sink (N=3)		4.4	7.12	87	1.62	5.89	109	3.51	3.67	76	0.08	1.7	4	-0.04	-0.05	17
Sink (N=4)		3.1	6.28	61	1.1	4.78	85	3.29	3.29	52	0.06	1.29	2	-0.01	0.07	13
Unimodal (N=2)		1.93	3.54	72	0.35	3.43	74	-0.01	-5.00	58	0.18	1.31	4	0	0.03	14
Unimodal (N=3)		1.72	3.19	43	0.31	3.15	54	0.13	-4.54	38	0.08	1.44	2	0	0.06	8
Unimodal (N=4)		1.27	2.8	30	0.24	2.77	42	0.18	-4.46	26	0.08	1.33	1	-0.02	-0.02	6
Crossmodal (N=2)		3.03	4.53	72	1.25	4.48	74	3.53	3.72	59	0.26	2.19	5	-0.02	0.01	15
Crossmodal (N=3)		2.15	3.7	44	1.01	4.11	55	3.3	3.15	38	0.2	1.6	2	0.01	0.06	9
Crossmodal (N=4)		1.45	3.02	31	0.63	3.57	43	3.04	2.86	26	0.07	1.25	1	0.02	0.04	7

Table 10. Causal tracing results obtained by patching hidden states after the MLP block.

Modality	Ablation	Qwen2.5-Omni (7B)			Qwen2.5-Omni (3B)			video-SALMONN-o1 (7B)			video-SALMONN2+ (7B)			video-SALMONN2+ (3B)		
		IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	#Tokens
Audio Dominant	All	4.68	2.65	1440	3.77	1.58	1440	5.80	9.33	1820	2.97	3.69	1210	0.84	1.42	1210
	Object	5.05	2.44	613	2.30	-0.03	580	2.66	5.18	852	2.31	3.02	500	0.45	0.93	448
	Random (N=2)	2.92	1.72	603	2.52	0.94	605	2.72	4.69	818	2.56	3.23	566	0.51	0.73	506
	Random (N=3)	2.55	1.14	363	2.28	0.48	355	0.93	2.44	514	2.19	3.01	360	0.39	0.50	298
	Random (N=4)	1.89	0.77	256	1.15	0.59	244	0.97	1.49	364	2.09	2.91	257	0.33	0.43	196
	Sink (N=2)	3.62	1.81	603	3.42	1.35	605	4.19	6.08	818	2.61	3.01	566	0.55	0.86	506
	Sink (N=3)	2.79	1.45	362	3.19	1.22	354	2.31	3.67	514	2.12	2.75	360	0.49	0.67	298
	Sink (N=4)	2.01	0.78	256	2.91	1.1	244	2.01	3.14	364	2.14	2.66	257	0.31	0.44	196
	Unimodal (N=2)	0.75	0.33	301	0.9	0.18	302	1.83	2.34	409	2.1	2.73	283	0.21	0.42	253
	Unimodal (N=3)	0.65	0.42	181	0.92	0.07	177	1.18	1.30	257	1.89	2.63	180	0.05	0.32	149
	Unimodal (N=4)	0.46	0.46	128	0.94	0.41	122	0.91	1.56	182	1.89	2.64	128	0.16	0.30	98
	Crossmodal (N=2)	3.89	1.87	301	3.16	1.31	303	2.04	3.52	409	2.24	2.84	283	0.54	0.78	253
	Crossmodal (N=3)	2.26	1.20	181	2.89	0.95	178	1.14	1.98	257	1.88	2.69	180	0.42	0.53	149
	Crossmodal (N=4)	1.96	0.9	128	2.74	0.99	122	1.05	1.67	182	1.84	2.58	129	0.37	0.48	98
	Video Dominant	All	10.66	13.29	249	4.89	8.66	250	-0.04	3.28	153	0.22	2.25	60	0.03	0.7
Object		7.49	8.21	149	3.45	6.14	149	-0.07	-0.49	78	0.14	1.47	7	0.01	0.13	7
Sink (N=2)		8.31	8.75	144	4.23	6.28	148	0.00	3.14	117	0.37	2.21	9	0.04	0.35	29
Sink (N=3)		7.69	7.51	87	3.56	5.45	109	0.00	2.80	76	0.21	1.62	4	0.02	0.24	17
Sink (N=4)		6.21	6.3	61	2.77	4.63	85	0.04	2.12	52	0.15	1.46	2	0.02	0.18	13
Random (N=2)		7.93	7.82	144	3.00	5.06	148	-0.06	0.62	117	0.31	1.49	9	0.02	0.42	29
Random (N=3)		6.21	5.4	87	2.75	4.22	109	-0.36	-1.94	76	0.17	1.42	4	0.03	0.25	17
Random (N=4)		4.62	3.98	61	2.36	3.55	85	-0.06	-2.81	52	0.10	1.34	2	0.03	0.25	13
Unimodal (N=2)		4.21	4.03	72	1.34	3.28	74	0.04	-4.87	58	0.17	1.55	4	0.02	0.33	14
Unimodal (N=3)		4.18	3.83	43	1.02	2.83	54	0.05	-4.70	38	0.16	1.33	2	0.03	0.23	8
Unimodal (N=4)		3.12	2.9	30	0.82	2.69	42	-0.02	2.66	59	0.06	1.32	1	0.02	0.27	6
Crossmodal (N=2)		5.86	5.16	72	3.17	4.32	74	-0.02	2.66	59	0.29	2.05	5	0.03	0.21	15
Crossmodal (N=3)		4.83	4.3	44	2.72	3.9	55	0.02	1.94	38	0.16	1.54	2	0.02	0.1	9
Crossmodal (N=4)		3.89	3.75	31	2.20	3.87	43	-0.06	1.10	26	0.14	1.39	1	0.02	0.08	7




Figure 10. Layer-wise causal patching results for the audio-dominant setting on Qwen2.5-Omni (7B). We patch all tokens from the non-dominant modality using a sliding window of 10 consecutive layers. cross-modal information exchange is most pronounced in the middle layers, indicating that these layers act as primary integration hubs.

Table 11. Causal tracing results on Qwen2.5-Omni(7B) under two alternative corruption methods: (1) injecting Gaussian noise into the raw input, and (2) replacing the encoder output embeddings with their mean.

Modality	Ablation	(1) Noise		(2) Encoder	
		IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑	IE <sub>clean</sub> ↑	IE <sub>corr</sub> ↑
Audio Dominant	All	9.6	5.3	6.7	2.5
	Object	5.1	2.4	4.1	0.5
	Random ( $N=2$ )	4.5	2.2	3.3	-0.3
	Random ( $N=3$ )	2.7	1.8	1.9	-1.5
	Sink ( $N=2$ )	6.2	2.9	3.6	-0.2
	Sink ( $N=3$ )	4.3	2.0	2.6	-0.9
	Unimodal ( $N=2$ )	0.7	0.2	0.6	-2.2
	Unimodal ( $N=3$ )	0.9	0.4	0.5	-2.5
	Crossmodal ( $N=2$ )	5.6	3.0	3.2	-0.5
	Crossmodal ( $N=3$ )	3.6	1.5	2.0	-1.6
Video Dominant	All	8.4	13.5	12.8	13.5
	Object	5.4	8.4	7.2	6.6
	Random ( $N=2$ )	4.5	6.4	7.8	7.0
	Random ( $N=3$ )	2.8	3.7	4.4	4.0
	Sink ( $N=2$ )	5.8	8.5	8.7	8.9
	Sink ( $N=3$ )	4.7	7.0	7.2	6.9
	Unimodal ( $N=2$ )	2.1	3.4	3.8	3.5
	Unimodal ( $N=3$ )	1.9	3.0	3.2	2.4
	Crossmodal ( $N=2$ )	3.3	4.4	5.2	4.3
	Crossmodal ( $N=3$ )	2.4	3.7	4.1	3.1

### B.1.6. MDS STATISTICS

We analyze MDS statistics of video and audio sink tokens for 100 samples (Tab. 12). Video and audio sink tokens exhibit positive and negative median values, respectively, indicating that sink tokens tend to receive more attention from their own modality than from complementary modality. However, the distribution is far from degenerate-given that MDS is bounded within  $[-1,1]$ , the sizable IQR and Std show considerable variation in the proportion of attention each sink token receives from two modalities, supporting Sec. 5.1 and Fig. 3.

Table 12. Distribution statistics of the Modality Dominance Score (MDS) for video and audio sink tokens, computed over 100 samples across five backbones.

Metric	Modality	Qwen2.5-Omni(7B)	Qwen2.5-Omni(3B)	video-SALMONN-o1(7B)	video-SALMONN2+(7B)	video-SALMONN2+(3B)
Median	Video	0.45	0.49	0.56	0.84	0.79
	Audio	-0.49	-0.50	-0.77	-0.18	-0.59
IQR	Video	0.34	0.34	0.39	0.13	0.22
	Audio	0.23	0.27	0.18	0.73	0.77
Std	Video	0.25	0.26	0.30	0.14	0.22
	Audio	0.26	0.27	0.16	0.55	0.46

## B.2. Object Hallucination Reduction

### B.2.1. OBJECT HALLUCINATION CASE ANALYSES

Tab. 13 quantifies the prevalence of such cross-modal disagreement cases and the hallucinated cases with audiovisual disagreement cases.

Table 13. Hallucination statistics on VGGSound-Animal dataset.

Metric	VGGSound-Animal
Total samples	363
Total audio-visual disagreement samples	272
Total hallucinated samples	175
Hallucinated samples with A-V disagreement	72

### B.2.2. ADDITIONAL RESULTS ON DIFFERENT MODELS

We extend our evaluation to verify whether ASD effectively mitigates hallucinations at the 3B-scale model. As shown in Tab. 14, applying ASD ( $\alpha = 0.3$ ,  $N = 2$ ) to Qwen2.5-Omni(3B) consistently reduces hallucination metrics across all three datasets, demonstrating the generalizability of our method.

We also note the distinct behavior of video-SALMONN2+. Unlike Qwen2.5-Omni and video-SALMONN-o1, the training paradigm of video-SALMONN2+ incorporates explicit mechanisms to suppress caption hallucinations. As a result, its baseline hallucination scores are already exceptionally low-showing an average  $C_s$  reduction of 43.9% and 21.1% across the three datasets for the 7B and 3B variants, respectively, compared to video-SALMONN-o1(7B) leaving a narrower margin for further improvement via training-free interventions.

Unlike approaches that require significant training overhead, ASD provides a **plug-and-play, training-free** alternative. Furthermore, we emphasize that our core contribution lies in the identification of *cross-modal information hubs*, which we leverage to effectively reduce captioning hallucinations across diverse models with varying scales and architectures.

### B.2.3. REVERSE ASD AS A COUNTERFACTUAL TEST

To further validate that the gains of ASD indeed stem from selectively promoting cross-modal sink tokens, we conduct a counterfactual experiment using a *reverse* variant of our method. Concretely, we reverse the signs in Eq. (6) and Eq. (7), so that the intervention decreases attention to cross-modal sink tokens while correspondingly increasing attention to unimodal sink tokens. We additionally scale the base guidance coefficient  $\alpha$  proportionally to the total sink attention attributed

Table 14. Quantitative evaluation of ASD applied to Qwen2.5-Omni(3B) using CHAIR for hallucination, and F1 scores for caption richness.

Dataset	Method	$C_s \downarrow$	$C_i \downarrow$	F1 $\uparrow$
VGGSound-Animal	Vanilla	51.24	41.00	52.88
	<b>ASD</b>	<b>43.25</b>	<b>38.52</b>	<b>50.44</b>
VGGSound-All	Vanilla	40.52	27.15	57.86
	<b>ASD</b>	<b>37.25</b>	<b>26.01</b>	<b>57.59</b>
Audioset	Vanilla	16.23	15.79	70.80
	<b>ASD</b>	<b>15.50</b>	<b>15.34</b>	<b>70.28</b>

to cross-modal sinks. This design ensures that when attention to cross-modal sinks is naturally high, the reverse ASD mechanism actively penalizes it while boosting attention to unimodal sinks.

Table 15. Effect of reverse ASD on Qwen2.5-Omni(7B) across three datasets.

Dataset	Method	$C_s \downarrow$	$C_i \downarrow$	F1 $\uparrow$
VGGSound-Animal	Vanilla	48.21	37.13	55.24
	Reverse ASD	46.83	41.74	49.14
VGGSound-All	Vanilla	30.70	20.67	58.69
	Reverse ASD	33.39	24.50	54.31
Audioset	Vanilla	8.92	10.93	69.73
	Reverse ASD	9.65	10.59	71.32

Tab. 15 presents the results of applying this reverse ASD on Qwen2.5-Omni(7B) across three datasets. As hypothesized, the overall performance deteriorates. However, we observe that the degradation is not perfectly symmetric to the improvements gained by the original ASD, consistent with the analysis in (Su et al., 2025). In hallucination-unprone states, the decoding process is firmly supported by stable evidence. Consequently, artificially increasing attention to unimodal sinks in these stable states does not easily induce new hallucinations. Furthermore, in already hallucination-prone states, where hallucinations are already occurring, further amplifying unimodal sink attention does not significantly induce additional hallucinations.

#### B.2.4. COMPARISON WITH INFERENCE-TIME INTERVENTIONS IN AVLLMS

We further compare ASD against two recent decoding methods designed for AVLLMs: AVCD (Jung et al., 2026) and FMD (Jung et al., 2025). As shown in Tab. 16, ASD consistently achieves the strongest overall performance across the three datasets, attaining the best  $C_s$  and  $C_i$  scores on every benchmark and the highest ALOHa on the two VGGSound subsets.

#### B.2.5. LATENCY OVERHEAD OF ASD

Although ASD proves effective, it comes at the cost of increased inference latency ( $3.7\times$  slower ms/token), arising from the sink identification step and the dual forward passes it requires.

### C. Qualitative Results

Fig. 11 shows qualitative results of our method. The baseline produces hallucinated object descriptions, whereas our method includes only genuine objects.

Table 16. Comparison with decoding-based hallucination mitigation baselines for AVLLMs on Qwen2.5-Omni(7B). We report ALOHa and CHAIR metrics for hallucination, and F1 for caption richness. **Bold** and underlined values indicate the best and second-best results within each dataset.

Dataset	Method	ALOHa $\uparrow$	$C_s$ $\downarrow$	$C_i$ $\downarrow$	F1 $\uparrow$
VGGSound-Animal	Vanilla	<u>40.7</u>	<u>48.2</u>	<u>37.1</u>	<b>55.2</b>
	AVCD	38.5	51.5	38.8	<u>53.7</u>
	FMD	37.3	50.9	40.2	52.9
	ASD	<b>42.7</b>	<b>36.9</b>	<b>34.1</b>	52.4
VGGSound-All	Vanilla	<u>35.0</u>	<u>30.7</u>	<b>20.6</b>	<b>58.6</b>
	AVCD	32.8	33.5	24.4	55.7
	FMD	32.5	37.9	25.9	<u>57.3</u>
	ASD	<b>38.8</b>	<b>29.6</b>	<u>21.7</u>	55.8
Audioset	Vanilla	38.2	<u>8.9</u>	<u>10.9</u>	69.7
	AVCD	<u>59.3</u>	10.8	12.1	69.2
	FMD	<b>63.7</b>	13.6	13.5	<u>71.2</u>
	ASD	38.2	<b>8.5</b>	<b>10.2</b>	<b>72.9</b>

### D. Limitations

Our work has several limitations. First, due to memory constraints, we were unable to evaluate our analysis on larger models such as Qwen3-Omni (30B)(Xu et al., 2025b), leaving the scalability of our findings to such scales as an open question. Second, our proposed ASD method is applicable only to captioning tasks and does not extend to general question answering. Finally, ASD introduces a non-trivial latency overhead, which may limit its practicality in latency-sensitive applications.



**Prompt :** Describe what you see and hear in a single sentence.

**Base**

In a serene outdoor setting, a **donkey** stands near a fence, its mouth open and tongue out, seemingly communicating or expressing itself audibly, while a **pig** oinks loudly in the background.

**Ours**

In this video, a **donkey** is standing in a fenced area, making a variety of funny noises, including oinking, screaming, and burping.



**Prompt :** Describe what you see and hear in a single sentence.

**Base**

A **lion** roars loudly, its mouth wide open, while a **pig** oinks in the background.

**Ours**

A **lion** is roaring loudly inside a fenced enclosure.

Figure 11. Qualitative results of our method. Text highlighted in red indicates hallucinated objects, while unhighlighted text denotes genuine objects.



Figure 11. Qualitative results of our method. Continued. Text highlighted in red indicates hallucinated objects, while unhighlighted text denotes genuine objects.