

RSA-Bench: Benchmarking Audio Large Models in Real-World Acoustic Scenarios

Anonymous ACL submission

Abstract

While Audio Large Models (ALLMs) have achieved remarkable proficiency, their robustness remains brittle in real-world deployment. Existing evaluations largely rely on synthetic Gaussian noise or simplistic single-source interference, failing to capture the intricate, multi-layered acoustic dynamics—or “Acoustic Ecology”—that characterize authentic physical environments. To bridge this ecological gap, we introduce **RSA-Bench**, a comprehensive robustness benchmark designed to stress-test ALLMs through high-fidelity auditory scene simulations. Unlike traditional methods, we construct evaluation samples by naturally superimposing diverse environmental soundscapes—spanning *Pasture*, *Extreme Weather*, *Classroom*, and *Outdoors*—onto clean speech signals across a spectrum of interference intensities. By evaluating models on six core tasks ranging from fundamental perception to complex reasoning, our study unveils three macro-level insights: **(I) The Perception-Cognition Gap:** Models maintain relative resilience in low-level recognition but suffer a **functional collapse** in high-order reasoning tasks under stress; **(II) Scenario Sensitivity:** “Vocal-like” interference (e.g., background laughter) proves significantly more destructive than mechanical noise, challenging the model’s auditory attention mechanisms; and **(III) The Denoising Paradox:** Standard speech enhancement often exacerbates performance degradation, as ALLMs prove highly sensitive to the semantic distortions introduced by denoising artifacts.

1 Introduction

In recent years, the intersection of Large Language Models (LLMs) and audio processing has given rise to Audio Large Models (ALLMs) (Goel et al., 2025; Yang et al., 2025b,a). By integrating audio encoders with pre-trained LLMs, these models have demonstrated remarkable capabilities across a wide range of tasks, including Automatic Speech

Recognition (ASR) (Ahlawat et al., 2025; Fatehifar et al., 2025; Liu et al., 2025b), speech translation (Sarim et al., 2025), and audio-based reasoning (Xie et al., 2025). Cutting-edge models have achieved impressive performance on standard benchmarks (Wang et al., 2025a; Kumar et al., 2025; Ma et al., 2025), exhibiting strong semantic understanding and instruction-following abilities when processing high-quality audio inputs.

However, the promising results obtained in controlled, noise-free environments often fail to translate to real-world deployment scenarios (Wang et al., 2025b; Atwany et al., 2025). Real-world acoustic environments are characterized by diverse, unavoidable background noises and multi-source interference. While previous works have established benchmarks for general audio capabilities (Yang et al., 2024; Wang et al., 2025a; Ahia et al., 2025), there is a systematic absence of evaluations that quantify how ALLMs behave under acoustic stress. Existing resources fail to reflect the complexity of a true “Acoustic Ecology,” (Wrightson, 2000; Pace et al., 2025) where target signals are inextricably intertwined with diverse background sounds. **Specifically, the magnitude of the performance gap between ideal and noisy conditions in ALLMs has not been sufficiently quantified, leaving the true robustness of these models in question.**

To address this fundamental limitation, we present *RSA-Bench*, a robustness benchmark designed to stress-test ALLMs within complex acoustic scenarios. Distinguished by its scale, the dataset covers more than 100,000 samples across six core tasks, ranging from basic ASR to high-order reasoning such as Math and QA. Specifically, the benchmark features a high-fidelity “Acoustic Ecology” constructed from four distinct environments: Pasture, Extreme Weather, Classroom, and Outdoor. To ensure realism, a multi-source superposition strategy is employed, naturally mixing 1 to 4 noise sources with the original signal. Ultimately, this

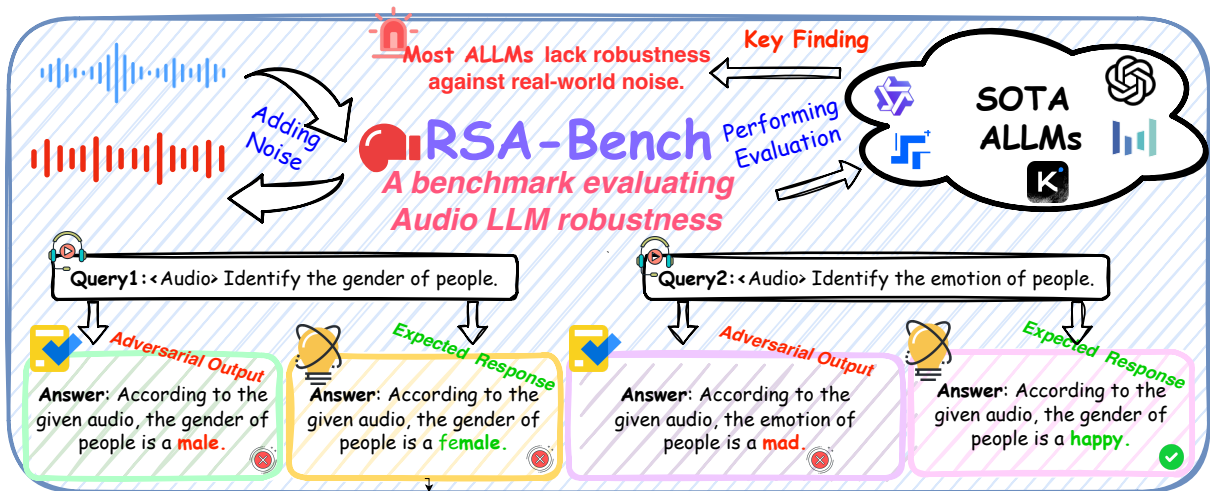


Figure 1: A framework of our RSA Benchmark for evaluating Audio-LLM robustness across six different tasks.

setting prioritizes ecological validity over artificial difficulty, simulating the complexity of real-world environments where target signals are inextricably intertwined with diverse background sounds.

As shown in Figure 1, our study reveals a stark contrast in model performance between clean and noisy inputs. We observe that most models exhibit a precipitous decline in capabilities as the acoustic environment becomes more complex, exposing a widespread vulnerability across current architectures. The degradation is particularly severe in tasks requiring precise semantic reasoning. Regarding mitigation, we applied three standard denoising methods to the noisy audio in an attempt to alleviate the negative impact of interference. However, we find that real-world noise proves to be remarkably persistent. Standard methods often struggle to effectively strip away this interference; instead, the attempt may disrupt the semantic integrity of the original audio, potentially leading to performance that is not only unrestored but further degraded.

Experimental Takeaways.

- **Widespread Robustness Vulnerability.** *RSA-benchmark* reveals a universal performance decline across diverse interference types, confirming that high capabilities in clean environments fail to translate to reliability in complex physical-world deployment.
- **The Perception-Cognition Gap.** Acoustic interference disproportionately impacts cognitive over perceptual capabilities. While models retain resilience in low-level tasks like gender recognition, they suffer a **functional collapse** in high-order reasoning under stress, exposing a critical bottleneck in complex semantic processing.

- **The Denoising Paradox.** External mitigation strategies often prove counterproductive. We find that standard speech enhancement algorithms frequently **exacerbate** errors, as ALLMs are significantly more sensitive to the spectral artifacts introduced by denoising than to the natural background noise itself.

2 Related Work

Audio Large Models. The landscape of audio processing has shifted dramatically from specialized models to general-purpose ALLMs (Zhang et al., 2023; Chu et al., 2024; Huang et al., 2024). Early works primarily focused on discriminative tasks such as ASR (Ahlawat et al., 2025; He and Whitehill, 2025). Recently, the integration of audio encoders with LLMs has empowered models like GPT-4o-Audio (Hurst et al., 2024) and Qwen2-Audio (Chu et al., 2024) to perform reasoning, instruction following, and multi-turn dialogue. The emergence of models like Qwen2.5-Omni (Xu et al., 2025a) further exemplifies the trend towards unified multimodal understanding (Zhang et al., 2025), where models process audio, text, and other modalities within a single end-to-end framework.

Robustness against Acoustic Interference. Robustness has been a longstanding pursuit in signal processing, traditionally measured by Word Error Rate (WER) in ASR systems under low Signal-to-Noise Ratio (SNR) (Song et al., 2025; Akomodi et al., 2025) conditions. In the era of ALLMs, the scope of robustness extends beyond recognition accuracy to encompass comprehensive understanding and reasoning capabilities in noisy contexts. Recent empirical studies have begun to explore the

negative impact of audio interference. For instance, recent work demonstrated that environmental noise can be utilized to bypass model safety mechanisms (Zhang and Lin, 2025; Peng et al., 2025; Chen et al., 2025), while other studies investigated how irrelevant audio acts as a distractor for text-based reasoning (Li et al., 2025a). However, the systematic impact of environmental noise on audio-centric cognitive tasks remains under-explored (Yang et al., 2024; Wang et al., 2025a). Furthermore, while speech enhancement (SE) (Yousif and Mahmmod, 2025; Jannu and Vanambathina, 2025; Huang et al., 2025) is a solution in traditional pipelines, its interaction with large-scale pre-trained encoders is complex. Our work provides a quantitative gap analysis and empirically examines the effectiveness of denoising methods. We find that applying off-the-shelf enhancement tools often fails to recover performance (Chondhekar et al., 2025), highlighting both the stubborn persistence of acoustic interference and the sensitivity of ALLMs to semantic distortions introduced by enhancement artifacts.

3 NoisyBench

Uniquely, *RSA-Bench* establishes the first framework to systematically investigate ALLM robustness against complex environmental noise. Instead of theoretical simulations, we construct four distinct real-world acoustic scenarios, each composed of representative audio elements designed to challenge specific aspects of model stability:

- **Pasture:** Represents an environment dominated by irregular animal vocalizations. We explicitly select non-stationary sounds from **cows**, **dogs**, **hens**, and **sheep** to test the model’s stability against sudden biological sounds.
- **Extreme Weather:** Simulates a complex acoustic environment with mixed interference types. This scenario combines continuous **heavy rain** and **wind** with sudden **thunderstorms** and tonal **wind chimes**, evaluating the model’s stability under varying acoustic pressure.
- **Classroom:** Replicates a typical indoor environment characterized by subtle but persistent human activity. We incorporate rhythmic **clock ticking** alongside sporadic human-generated noises such as **coughing**, **keyboard typing**, and **drinking**, simulating a scenario where background activities compete with the target speech.
- **Outdoors:** Represents an open-air environment.

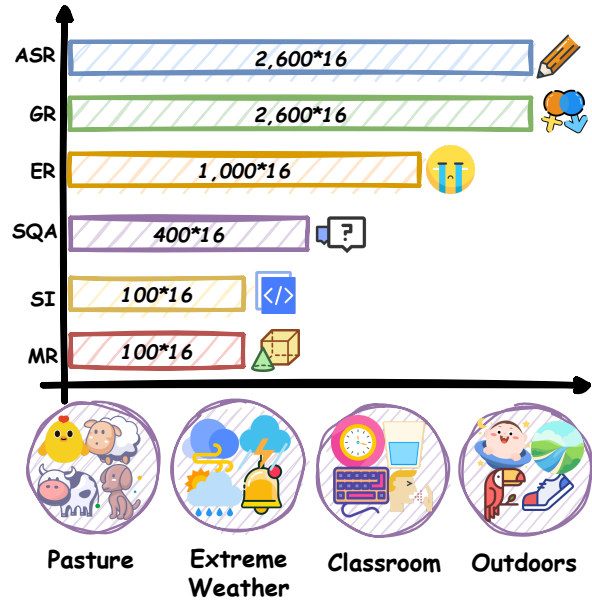


Figure 2: Overview of the RSA-Bench data composition, which covers 6 tasks and totals over 100,000 samples.

To ensure ecological fidelity, we synthesize a soundscape featuring **children playing**, **bird chirping**, **flowing streams**, and texture-specific **footsteps on grass**. This tests the model’s adaptability to unstructured acoustic events.

3.1 Data Construction

To systematically evaluate model robustness, we construct four distinct variations for each of the four predefined acoustic scenarios by varying the number of superimposed real-world interference, specifically setting $K \in \{1, 2, 3, 4\}$ to represent increasing levels of environmental complexity. For each individual audio sample in the dataset, this construction process follows four sequential steps: source collection, temporal alignment, energy alignment, and superposition.

Step 1. Source Collection. The construction of NoisyBench begins with the curation of high-quality source materials. We aggregate data from two distinct sources:

- **Clean Audio Stream:** We select samples from six representative datasets covering both perception tasks and reasoning tasks.
- **Noise Audio Stream:** To simulate authentic acoustic ecologies, we utilize recordings from the Environmental Sound Classification (ESC) (Piczak, 2015) subset of DynamicSuperb. These are manually categorized into four distinct scenarios: *Pasture*, *Extreme Weather*, *Classroom*,

and *Outdoor*.

Step 2. Temporal Alignment. Upon obtaining the source materials, we define the discrete-time clean audio signal $s[n]$ of length N . For each sample, we randomly select K noise clips from a target environmental category ($K \in \{1, \dots, 4\}$). Let $w_k[n]$ represent the k -th raw noise signal of length M_k . To address the duration mismatch between the clean audio and the noise, we apply a temporal alignment operator. We generate the aligned noise sequence $\tilde{w}_k[n]$ using a modulo operation:

$$\tilde{w}_k[n] = w_k[n \bmod M_k], \quad \text{for } 0 \leq n < N. \quad (1)$$

This formulation unifies two behaviors: if the noise is shorter than the clean audio ($M_k < N$), it is cyclically tiled to fill the duration; if the noise is longer ($M_k > N$), it is automatically truncated to match the target length N . This ensures continuous background coverage.

Step 3. RMS-based Energy Alignment. To establish a consistent interference intensity, we normalize the energy of the noise audio to strictly match that of the clean audio. We first calculate the Root Mean Square (RMS) energy for the clean audio (R_s) and the aligned noise audio (R_{w_k}):

$$R_s = \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} s^2[n]}, \quad R_{w_k} = \sqrt{\frac{1}{N} \sum_{n=0}^{N-1} \tilde{w}_k^2[n]}. \quad (2)$$

We then derive an adaptive scaling factor λ_k to align the noise energy to the speech energy:

$$\lambda_k = \frac{R_s}{R_{w_k}}. \quad (3)$$

In our experiments, we fix $\lambda_k = 1$.

Step 4. Superposition and Dynamic Constraint. Finally, we generate the noisy speech sample by linearly superimposing the clean audio and the scaled real-world interference. To ensure the audio data remains within the valid amplitude range, we apply a hard constraint function. The final evaluation sample $x[n]$ is formulated as:

$$x[n] = \text{clip} \left(s[n] + \sum_{k=1}^K (\tilde{w}_k[n] \cdot \lambda_k), -1, 1 \right), \quad (4)$$

where $\text{clip}(v, -1, 1)$ restricts the amplitude values to the interval $[-1, 1]$.

By executing the aforementioned four steps for each clean sample across all scenarios and the four interference levels ($K = 1$ to 4), this combinatorial design results in 4 scenarios \times 4 intensity levels = 16 unique configurations per original sample. Together with the original clean version, NoisyBench provides a total of 17 test conditions per sample, enabling a fine-grained analysis of model robustness as environmental complexity scales.

3.2 Task Taxonomy and Definitions

To comprehensively disentangle the impact of environmental complexity on different model capabilities, we categorize the six evaluation tasks into two distinct **categories: Perception & Paralinguistics** and **Cognitive Reasoning**.

3.2.1 Perception & Paralinguistics.

These tasks assess the model’s fundamental ability to perceive acoustic signals and extract specific attributes, evaluating whether the model can maintain signal fidelity under environmental interference.

ASR. ASR aims to transcribe spoken content into verbatim text. This task measures the model’s robustness in preserving linguistic information against environmental masking. We use source samples from LibriSpeech (Panayotov et al., 2015) to evaluate phonetic recognition accuracy under complex acoustic conditions.

Gender Recognition (GR). This task evaluates the ability to discern speaker identity traits based on vocal characteristics. Established upon the IEMO-CAP dataset (Busso et al., 2008), it challenges the model to isolate the speaker’s biological features from background environments, testing the robustness of acoustic feature extraction.

Emotion Recognition (ER). Emotion is a critical paralinguistic element conveyed through prosody and tone. Utilizing MELD (Porcia et al., 2019) as the source, this task requires the model to interpret the speaker’s emotional state.

3.2.2 Cognitive Reasoning.

These tasks require the model to perform logical processing based on the audio inputs. They test the robustness of ALLMs’ cognitive capabilities against acoustic interference.

Mathematical Reasoning (MR). This task involves extracting numerical values to perform calculations. We utilize SpokenMQA (Wei et al., 2025) to evaluate the reliability of the model’s reasoning process under acoustic stress. This task

K	Models										
	Qwen2	SALMONN	SeaLLMs	Phi	MERaLION	Step2	MiniCPM	Qwen-Turbo	Qwen2.5-Omni	Qwen3-Omni	gpt-mini
ASR (WER ↓)											
$K=0$	3.45	10.49	5.52	1.67	2.34	3.90	2.95	23.78	23.32	1.72	50.01
$K=1$	8.49 \downarrow 5.04	24.47 \uparrow 13.98	25.49 \uparrow 19.97	7.07 \uparrow 5.40	11.63 \uparrow 9.29	7.59 \uparrow 3.69	21.08 \uparrow 18.13	27.95 \uparrow 4.17	28.89 \uparrow 5.57	5.70 \uparrow 3.98	64.69 \uparrow 14.68
$K=2$	19.67 \uparrow 16.22	124.79 \uparrow 114.3	51.13 \uparrow 45.61	19.54 \uparrow 17.87	30.89 \uparrow 28.55	20.47 \uparrow 16.57	57.34 \uparrow 54.39	42.30 \uparrow 38.52	45.18 \uparrow 41.86	48.41 \uparrow 46.69	86.97 \uparrow 76.96
$K=3$	35.97 \uparrow 32.52	317.33 \uparrow 306.8	125.50 \uparrow 119.9	42.89 \uparrow 41.22	55.35 \uparrow 53.01	34.49 \uparrow 30.59	89.09 \uparrow 86.14	61.54 \uparrow 57.76	61.57 \uparrow 58.25	259.56 \uparrow 257.8	107.27 \uparrow 97.26
$K=4$	54.75 \uparrow 51.30	509.11 \uparrow 498.6	279.27 \uparrow 273.7	81.12 \uparrow 79.45	76.04 \uparrow 73.70	66.67 \uparrow 62.77	121.17 \uparrow 118.2	96.42 \uparrow 92.64	93.27 \uparrow 89.95	557.20 \uparrow 555.5	118.39 \uparrow 108.38
ER (Score ↑)											
$K=0$	51.53	40.54	47.81	49.92	52.61	56.82	55.33	52.99	52.91	47.20	30.61
$K=1$	35.33 \downarrow 16.20	30.87 \downarrow 9.67	22.45 \downarrow 25.36	24.13 \downarrow 25.79	56.59 \uparrow 3.98	39.27 \downarrow 17.55	30.03 \downarrow 25.30	22.91 \downarrow 30.08	23.18 \downarrow 29.73	34.10 \downarrow 13.10	5.67 \downarrow 24.94
$K=2$	33.68 \downarrow 17.85	30.46 \downarrow 10.08	20.99 \downarrow 26.82	23.71 \downarrow 26.21	55.82 \uparrow 3.21	38.58 \downarrow 18.24	29.34 \downarrow 25.99	25.02 \downarrow 27.97	24.90 \downarrow 28.01	33.56 \downarrow 13.64	8.93 \downarrow 21.68
$K=3$	32.84 \downarrow 18.69	30.46 \downarrow 10.08	21.22 \downarrow 26.59	24.09 \downarrow 25.83	56.51 \uparrow 3.90	38.39 \downarrow 18.43	29.84 \downarrow 25.49	14.10 \downarrow 38.89	13.60 \downarrow 39.31	33.41 \downarrow 13.79	1.07 \downarrow 29.54
$K=4$	30.84 \downarrow 20.69	30.23 \downarrow 10.31	20.22 \downarrow 27.59	23.94 \downarrow 25.98	55.05 \uparrow 2.44	39.69 \downarrow 17.13	29.42 \downarrow 25.91	10.57 \downarrow 42.42	12.34 \downarrow 40.57	34.10 \downarrow 13.10	0.23 \downarrow 30.38
GR (Score ↑)											
$K=0$	96.02	82.37	79.87	38.65	85.26	86.95	93.43	91.63	91.53	95.92	-
$K=1$	93.62 \downarrow 2.40	82.67 \uparrow 0.30	71.04 \downarrow 8.83	38.65 -	82.97 \downarrow 2.29	85.96 \downarrow 0.99	91.33 \downarrow 2.10	91.53 \downarrow 0.10	91.04 \downarrow 0.49	92.53 \downarrow 3.39	-
$K=2$	90.03 \downarrow 5.99	82.66 \uparrow 0.29	74.19 \downarrow 5.68	32.07 \downarrow 6.58	84.06 \downarrow 1.20	81.67 \downarrow 5.28	90.13 \downarrow 3.30	89.84 \downarrow 1.79	90.14 \downarrow 1.39	91.33 \downarrow 4.59	-
$K=3$	87.64 \downarrow 8.38	65.83 \downarrow 16.54	73.23 \downarrow 6.64	27.69 \downarrow 10.96	81.77 \downarrow 3.49	83.76 \downarrow 3.19	85.16 \downarrow 8.27	88.55 \downarrow 3.08	89.54 \downarrow 1.99	88.35 \downarrow 7.57	-
$K=4$	81.77 \downarrow 14.25	59.66 \downarrow 22.71	75.66 \downarrow 4.21	21.41 \downarrow 17.24	76.49 \downarrow 8.77	77.19 \downarrow 9.76	81.87 \downarrow 11.56	86.35 \downarrow 5.28	87.05 \downarrow 4.48	88.94 \downarrow 6.98	-
MR (Acc ↑)											
$K=0$	66.00	18.00	62.00	3.00	74.00	75.00	75.00	88.00	89.00	91.00	93.00
$K=1$	43.00 \downarrow 23	5.00 \downarrow 13	29.00 \downarrow 33	1.00 \downarrow 2	46.00 \downarrow 28	47.00 \downarrow 28	42.00 \downarrow 33	48.00 \downarrow 40	39.00 \downarrow 50	63.00 \downarrow 28	49.00 \downarrow 44
$K=2$	23.00 \downarrow 43	2.00 \downarrow 16	1.00 \downarrow 52	2.00 \downarrow 1	27.00 \downarrow 47	26.00 \downarrow 49	18.00 \downarrow 57	18.00 \downarrow 70	16.00 \downarrow 73	40.00 \downarrow 51	16.00 \downarrow 77
$K=3$	6.00 \downarrow 60	0.00 \downarrow 18	1.00 \downarrow 61	2.00 \downarrow 1	7.00 \downarrow 67	12.00 \downarrow 63	7.00 \downarrow 68	7.00 \downarrow 81	4.00 \downarrow 85	18.00 \downarrow 73	6.00 \downarrow 87
$K=4$	4.00 \downarrow 62	0.00 \downarrow 18	0.00 \downarrow 62	1.00 \downarrow 2	3.00 \downarrow 71	6.00 \downarrow 69	1.00 \downarrow 74	4.00 \downarrow 84	4.00 \downarrow 85	5.00 \downarrow 86	3.00 \downarrow 90
SQA (Score ↑)											
$K=0$	79.85	79.90	78.58	85.74	80.69	81.37	82.94	82.50	83.82	80.98	86.62
$K=1$	77.21 \downarrow 2.64	73.14 \downarrow 6.76	73.92 \downarrow 4.66	86.42 \uparrow 0.68	80.29 \downarrow 4.40	79.31 \downarrow 2.06	82.45 \downarrow 0.49	82.65 \uparrow 0.15	81.37 \downarrow 2.45	80.83 \downarrow 0.15	86.18 \downarrow 0.44
$K=2$	73.97 \downarrow 5.88	67.84 \downarrow 12.06	65.93 \downarrow 12.65	80.69 \downarrow 5.05	77.60 \downarrow 3.09	76.47 \downarrow 4.90	80.00 \downarrow 2.94	79.71 \downarrow 2.79	78.97 \downarrow 4.85	81.03 \uparrow 0.05	86.18 \downarrow 0.44
$K=3$	67.21 \downarrow 12.64	63.82 \downarrow 16.08	58.38 \downarrow 20.20	80.05 \downarrow 5.69	73.68 \downarrow 7.01	69.07 \downarrow 12.30	75.78 \downarrow 7.16	70.20 \downarrow 12.30	73.53 \downarrow 10.29	75.74 \downarrow 5.24	83.38 \downarrow 3.24
$K=4$	62.25 \downarrow 17.60	62.74 \downarrow 17.16	55.74 \downarrow 22.84	73.28 \downarrow 12.46	71.08 \downarrow 9.61	61.27 \downarrow 20.10	71.37 \downarrow 11.57	66.62 \downarrow 15.88	66.96 \downarrow 16.86	73.53 \downarrow 7.45	79.17 \downarrow 7.45
SI (Score ↑)											
$K=0$	49.60	58.40	62.00	33.20	71.00	58.20	72.40	78.20	76.60	82.60	78.20
$K=1$	43.20 \downarrow 6.4	51.60 \downarrow 6.8	39.20 \downarrow 22.8	26.40 \downarrow 6.8	67.40 \downarrow 3.6	51.60 \downarrow 6.6	59.80 \downarrow 12.6	70.40 \downarrow 7.8	68.60 \downarrow 8.0	70.20 \downarrow 12.4	76.00 \downarrow 2.2
$K=2$	32.60 \downarrow 17.0	55.20 \downarrow 3.2	21.60 \downarrow 40.4	18.40 \downarrow 14.8	55.20 \downarrow 15.8	36.40 \downarrow 21.8	46.80 \downarrow 25.6	53.80 \downarrow 24.4	56.60 \downarrow 20.0	58.40 \downarrow 24.2	52.00 \downarrow 26.2
$K=3$	19.80 \downarrow 29.8	57.40 \downarrow 1.0	10.00 \downarrow 52.0	17.60 \downarrow 15.6	37.40 \downarrow 33.6	23.80 \downarrow 34.4	30.20 \downarrow 42.2	36.20 \downarrow 42.0	34.20 \downarrow 42.4	41.40 \downarrow 41.2	29.60 \downarrow 48.6
$K=4$	6.60 \downarrow 43.0	54.60 \downarrow 3.8	3.60 \downarrow 58.4	11.80 \downarrow 21.4	17.20 \downarrow 53.8	9.80 \downarrow 48.4	12.00 \downarrow 60.4	21.00 \downarrow 57.2	20.00 \downarrow 56.6	17.80 \downarrow 64.8	6.40 \downarrow 71.8

Table 1: Results under varying noise-source count ($K=0-4$) in the *Outdoors* acoustic scenario. Variations relative to $K=0$ are indicated with \uparrow (increase) and \downarrow (decrease). Best (worst) results within each K row are shown in **bold** (underlined). All data values are presented with the unit of percent (%).

assesses whether the model can accurately interpret numerical information and perform correct calculations despite environmental distractions.

Speech Question Answering (SQA). Simulating real-world comprehension requires the model to understand spoken passages and answer logic-dependent questions. Based on SLUE Phase-2 (Shon et al., 2023), it tests the model’s ability to retrieve specific facts and perform deductive reasoning when the context is perturbed.

Speech Instruction Following (SI). Mirroring natural human-computer interaction, this task evaluates whether the model can understand and execute complex, open-ended instructions delivered via audio. Using the OpenHermes (Shon et al., 2023) instruction set, we assess the model’s capability to parse user intent and adhere to complex constraints within a realistic acoustic environment.

4 Experiments

Methods. We select a diverse set of representative ALLMs, ranging from unified proprietary models to open-source frameworks, including Qwen2-Audio-7B-Instruct (Chu et al., 2024), Qwen2.5-Omni-7B (Xu et al., 2025a), SeaLLMs-Audio-7B (Liu et al., 2025a), MERaLION-AudioLLM-Whisper-SEA-LION (He et al., 2024), phi-4-multimodal-instruct (Abouelenin et al., 2025), Step-Audio-2-mini (Wu et al., 2025), SALMONN-7B

(Tang et al., 2023), MiniCPM-o-2.6 (Yao et al., 2024), Qwen3-Omni-Flash (Xu et al., 2025b), Qwen-Omni-Turbo (Tang et al., 2023), and GPT-4o-mini-audio (Achiam et al., 2023). These models cover various architectures and training strategies, providing a comprehensive view of the current landscape.

Evaluation Metrics. We adopt task-specific metrics to ensure rigorous assessment. For ASR, we employ WER, where lower values indicate better robustness. For Mathematical Reasoning (MR), we calculate Accuracy (Acc) based on exact numerical matching. For tasks involving semantic comprehension or paralinguistic classification (SQA, SI, ER, GR), we utilize an LLM-as-a-Judge (Zheng et al., 2023; Li et al., 2025b) approach. Specifically, GPT-4o-mini serves as the evaluator, scoring model responses against ground truths on a 0-100 scale, focusing on semantic correctness and instruction compliance. Additionally, we evaluate all models on a Clean Baseline (original, uncorrupted audio) to quantify the relative performance degradation under noisy conditions.

Inference Settings. We evaluate each ALLM across the 17 distinct acoustic conditions per sample as defined in Sec. 3.1. This includes the original Clean Baseline and the 16 noisy configurations spanning the four real-world scenarios and complexity levels ($K = 1$ to 4). This benchmarking across a predefined stress gradient allows us to

379 quantify the performance gap between ideal and
380 complex environments.

381 5 Main Results

382 In this section, we analyze the performance of 11
383 ALLMs under various real-world interference con-
384 ditions. Our analysis explores the limits of model
385 robustness through four Research Questions (RQs)
386 covering the impact of environmental complexity,
387 task-specific sensitivity, acoustic ecologies, and
388 model-wise performance patterns.

389 5.1 Impact of Real-world Interference (RQ1)

390 **Obs 1: Differentiation within Perception Tasks.**
391 In Table 1, we observe distinct behaviors across
392 perception tasks under noise. **GR** remains highly
393 robust; for instance, *Qwen3-Omni* retains high ac-
394 curacy even at $K=4$ (95.92% \rightarrow 88.94%). In con-
395 trast, **ER** proves fragile: *Qwen-Turbo* plunges from
396 52.99% (Clean) to 10.57% ($K=4$), while *gpt-mini*
397 collapses completely (30.61% \rightarrow 0.23%). This sug-
398 gests coarse-grained features (gender) withstand
399 interference, whereas nuanced affective cues are
400 easily distorted.

401 **Obs 2: Vulnerability of Reasoning-Heavy Tasks.**
402 Tasks demanding complex cognition, such as **MR**,
403 degrade much faster than perception tasks. *Step2*
404 accuracy drops sharply from 75.00% ($K=0$) to
405 47.00% ($K=1$) and finally 6.00% ($K=4$). This
406 steep downward trajectory is consistent across mod-
407 els, with most falling below 5% accuracy under
408 maximum noise. Such rapid failure indicates that
409 noise severely disrupts the precise information ex-
410 traction required for multi-step reasoning.

411 5.2 Disparity Across Different Tasks (RQ2)

412 **Obs 3: ASR Functional Collapse.** Models show
413 robustness at low interference ($K = 1$), with *phi-4*
414 maintaining ~ 0.07 WER. However, performance
415 degrades sharply at $K = 4$, leading to “functional
416 collapse” where models ignore prompts in favor
417 of conversational output or gibberish. Notably,
418 *Qwen3-omni-flash* surges from **0.06** WER ($K = 1$)
419 to **5.57** in *Outdoors-4*.

420 **Obs 4: Robustness of Gender Recognition.** **GR**
421 demonstrates exceptional resilience. Even in ex-
422 treme *Outdoors* ($K = 4$) conditions, accuracy
423 remains $\sim 80\%$ (dropping slightly from **84.2%**
424 clean). This indicates biological vocal features are
425 far more stable against acoustic interference than
426 semantic features.

Obs 5: Erosion of Semantic Understanding. 427
Conversely, tasks requiring complex semantic pro- 428
cessing (**SI**, **SQA**) suffer severe degradation. Aver- 429
age **SI** scores plummet from **65.5** (Clean) to **16.4** 430
at $K = 4$. Noise masking critical tokens disrupts 431
intent recognition, causing models to fail in instruc- 432
tion adherence. 433

434 5.3 Scenario-specific Impact (RQ3)

Table 2 compares model performance at $K = 3$. 435

Obs 6: Extreme Challenge of Outdoors. The 436
Outdoors scenario consistently imposes the heav- 437
iest toll. Analysis suggests that “vocal-like” non- 438
verbal sounds (e.g., laughter) overlap significantly 439
with human speech frequencies. Consequently, 440
Qwen3-Omni’s ASR WER surges to **259.56** at 441
 $K = 3$ —a massive gap compared to **4.77** in the 442
Classroom—indicating models struggle to filter in- 443
terference that mimics target speech traits. 444

Obs 7: Resilience in Classroom Scenario. Con- 445
versely, models perform best in the **Classroom**, 446
where *Qwen3-Omni* reaches a peak MR score of 447
72.00 at $K = 3$. The discrete, rhythmic nature of 448
noises (e.g., typing) creates “acoustic gaps,” allow- 449
ing ALLMs to capture speech information during 450
silent intervals to effectively mitigate interference. 451

Obs 8: Spectral Masking in Extreme Weather. 452
Extreme Weather degrades performance through 453
continuous broadband noise (e.g., rain). Acting 454
as a “spectral blanket,” this interference uniformly 455
blurs acoustic details, making fine-grained phonetic 456
distinction significantly harder than in sparse noise 457
environments and causing steady WER increases. 458

459 5.4 Cross-Model Capability (RQ4)

Obs 9: Variability in Model Resistance. 460
ALLMs exhibit significant performance gaps un- 461
der pressure. In **ASR**, *qwen2_audio_7b_instruct* 462
and *step_audio_2_mini* show resilience, maintain- 463
ing low WERs (~ 0.35) even at $K = 3$. In **SI**, 464
SALMONN-7B sustains a score of **54.60** under ex- 465
treme interference ($K = 4$), whereas *SeaLLMs* 466
drops drastically from **0.62** ($K = 0$) to **0.036**. 467
These results highlight distinct robustness bound- 468
aries across architectures. 469

470 6 Robustness Mitigation via Denoising

Interference in real-world acoustic environments 471
can significantly degrade the performance of 472
ALLMs. To bridge this gap, we conduct a series of 473

Scenario	Models										
	Qwen2	SALMONN	SeaLLMs	Phi	MERaLION	Step2	MiniCPM	Qwen-Turbo	Qwen2.5-Omni	Qwen3-Omni	gpt-mini
<i>ASR (WER ↓)</i>											
Pasture	14.48	73.41	36.55	12.80	22.49	14.10	38.56	40.69	43.17	10.61	87.32
Extreme Weather	24.64	170.19	64.75	34.08	38.33	25.68	64.42	46.51	49.77	43.18	<u>85.35</u>
Classroom	9.95	27.14	27.63	8.81	10.34	8.66	18.39	26.32	26.06	4.77	<u>66.27</u>
Outdoors	35.97	<u>317.33</u>	125.50	42.89	55.35	34.49	89.09	61.54	61.57	259.56	107.27
<i>ER (Score ↑)</i>											
Pasture	29.69	25.13	20.46	21.53	50.46	34.94	26.67	19.96	19.50	28.47	<u>1.88</u>
Extreme Weather	33.91	29.81	22.18	24.41	54.60	39.54	30.38	17.82	17.66	35.40	<u>3.83</u>
Classroom	34.98	29.77	22.07	23.37	55.36	38.89	27.84	22.99	23.37	32.80	<u>4.71</u>
Outdoors	32.84	30.46	21.22	24.09	56.51	38.39	29.84	14.10	13.60	33.41	<u>1.07</u>
<i>GR (Score ↑)</i>											
Pasture	95.12	76.29	69.62	<u>31.47</u>	82.47	83.76	87.75	92.43	91.43	95.22	–
Extreme Weather	92.93	69.22	69.42	<u>30.08</u>	84.16	83.76	90.84	92.83	92.33	95.22	–
Classroom	95.52	72.81	77.06	<u>33.96</u>	84.06	89.74	85.86	94.32	95.22	95.92	–
Outdoors	87.65	65.84	73.24	<u>27.69</u>	81.77	83.76	85.16	88.55	89.54	88.35	–
<i>MR (Acc ↑)</i>											
Pasture	20.00	<u>3.00</u>	22.00	1.00	28.00	30.00	19.00	20.00	23.00	44.00	29.00
Extreme Weather	15.00	<u>0.00</u>	8.00	1.00	15.00	14.00	10.00	13.00	10.00	29.00	24.00
Classroom	33.00	<u>4.00</u>	35.00	3.00	52.00	50.00	48.00	47.00	46.00	72.00	51.00
Outdoors	6.00	<u>0.00</u>	1.00	2.00	7.00	12.00	7.00	7.00	4.00	18.00	6.00
<i>SQA (Score ↑)</i>											
Pasture	75.25	<u>68.92</u>	70.59	82.79	73.68	75.69	79.75	79.31	79.02	81.76	86.32
Extreme Weather	71.47	68.04	<u>64.22</u>	80.54	76.62	74.51	76.52	76.32	80.10	81.47	84.85
Classroom	75.54	<u>72.94</u>	75.88	83.68	81.76	77.21	83.58	83.97	83.68	82.30	86.13
Outdoors	67.21	63.82	<u>55.74</u>	73.28	73.68	69.07	71.37	70.20	73.53	75.74	83.38
<i>SI (Score ↑)</i>											
Pasture	35.20	53.20	26.60	<u>22.40</u>	55.00	40.80	58.60	56.00	53.80	62.20	59.40
Extreme Weather	29.00	54.20	<u>23.40</u>	<u>20.60</u>	45.60	37.80	39.00	49.80	48.40	57.40	51.80
Classroom	35.00	54.80	<u>40.20</u>	<u>20.40</u>	65.80	40.00	67.80	69.40	72.80	77.00	75.00
Outdoors	19.80	57.40	<u>10.00</u>	17.60	37.40	23.80	30.20	36.20	34.20	41.40	29.60

Table 2: Results under a fixed noise-source count ($K=3$) across four acoustic scenarios. Each block groups rows by task; values are reported as percentages. For ASR we report word error rate (WER, lower is better), while other tasks use accuracy or LLM-judge scores (higher is better). Arrows indicate the preferred direction. Best (worst) results within each scenario row are shown in **bold** (underlined).

experiments utilizing various denoising algorithms to evaluate whether current speech enhancement techniques can effectively restore model performance. We select 4 representative denoising algorithms for evaluations (detailed implementations are provided in Appendix A):

- **noisereduce (Sainburg et al., 2020)**: A traditional stationary noise reduction method based on spectral gating.
- **RNNoise (Valin, 2018)**: A hybrid approach combining classic signal processing with Recurrent Neural Networks (RNNs).
- **Audio-Denoising (Ali and Shemi, 2015)**: A Wavelet Transform approach.
- **DeepFilterNet (Schröter et al., 2023)**: A low-latency speech enhancement framework utilizing complex deep filtering.

We apply these methods to clean the noisy samples before feeding them into the ALLMs. To evaluate whether model performance can be further improved, we select 3 top-performing models, including Qwen2-Audio-7B-Instruct, MERaLiON-AudioLLM, and Step-Audio-2-mini. The evaluation is conducted on two contrasting scenarios: **Pasture** and **Classroom**. The performance of the enhanced audio and conduct a parallel comparison with the noisy baseline and clean reference data.

Obs 10: Performance Decline Following Denoising. As evidenced in Table 5, we observe a counter-intuitive trend: applying external denoising algorithms prior to inference frequently degrades rather than enhances model performance. This suggests ALLMs, likely exposed to diverse environmental sounds during pre-training, are far more robust to natural background noise than to signal distortion and spectral artifacts from enhancement techniques. For instance, in the **ASR** task with *Qwen2* ($K=1$), the Word Error Rate deteriorates from a baseline of 4.24 to 5.62 with *audio_denoising*, and worsens dramatically to 12.90 with *noisereduce*. A similar regression is seen in the **SI** semantic task, where *DeepFilterNet* reduces accuracy from 47.20 to 43.00. These results indicate aggressive filtering inadvertently compromises critical acoustic cues (phonemic boundaries, speaker timbre), and the negative impact of these artifacts outweighs the theoretical benefits of noise reduction.

Obs 11: Horizontal Comparison of Denoising Methodologies. A horizontal comparison of different techniques reveals that traditional signal processing methods are far more destructive to ALLM performance than modern deep learning approaches. For instance, in the $K = 4$ ASR task (*Classroom*), *noisereduce* causes the WER

(a) ASR (WER ↓)						(b) ER (Score ↑)					
Method	$K = 0$	$K = 1$	$K = 2$	$K = 3$	$K = 4$	Method	$K = 0$	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise		4.24	6.47	9.95	14.63	Noise		37.47	36.28	34.98	33.72
NoiseReduce		<u>12.90</u>	<u>24.16</u>	<u>38.56</u>	<u>55.74</u>	NoiseReduce		<u>32.11</u>	<u>29.16</u>	35.21	25.79
AudioDenoise	3.45	5.62	10.88	18.67	30.61	AudioDenoise	51.53	36.28	34.14	<u>34.21</u>	<u>23.64</u>
PyRNNoise		7.73	15.08	24.62	36.48	PyRNNoise		32.80	29.35	35.75	31.11
DeepFilterNet		7.71	14.03	22.51	32.62	DeepFilterNet		34.56	32.34	34.83	32.64

(c) GR (Score ↑)						(d) MR (Acc ↑)					
Method	$K = 0$	$K = 1$	$K = 2$	$K = 3$	$K = 4$	Method	$K = 0$	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise		95.22	95.62	95.52	94.62	Noise		62.00	46.00	33.00	26.00
NoiseReduce		88.55	84.56	83.47	79.68	NoiseReduce		29.00	18.00	9.00	4.00
AudioDenoise	96.02	94.82	92.73	93.53	93.23	AudioDenoise	66.00	55.00	37.00	27.00	14.00
PyRNNoise		92.63	91.43	92.13	90.94	PyRNNoise		46.00	33.00	21.00	16.00
DeepFilterNet		92.23	92.83	93.43	92.23	DeepFilterNet		44.00	32.00	21.00	18.00

(e) SQA (Score ↑)						(f) SI (Score ↑)					
Method	$K = 0$	$K = 1$	$K = 2$	$K = 3$	$K = 4$	Method	$K = 0$	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise		78.68	78.28	75.54	72.45	Noise		47.20	45.60	35.00	29.60
NoiseReduce		<u>76.13</u>	<u>73.48</u>	<u>70.20</u>	<u>63.77</u>	NoiseReduce		<u>33.20</u>	<u>29.20</u>	<u>21.60</u>	<u>8.40</u>
AudioDenoise	79.85	78.87	75.20	70.74	65.20	AudioDenoise	49.60	42.40	47.40	33.20	22.60
PyRNNoise		79.71	76.08	74.80	70.49	PyRNNoise		47.60	40.20	35.80	25.40
DeepFilterNet		81.32	77.01	73.53	72.70	DeepFilterNet		43.00	41.60	35.60	27.80

Table 3: Denoising ablation in the Classroom scenario for Qwen2. Rows are denoising methods (Noise = no denoise); columns are noise-source count $K = 0 \dots 4$. Values are percentages (numbers < 1 are multiplied by 100). For ASR we report WER (lower is better); other tasks are higher-is-better. Best (worst) per K column is in **bold** (underlined).

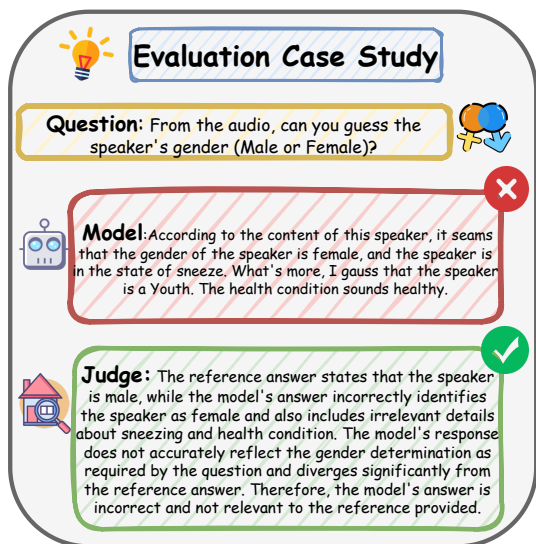


Figure 3: A framework of our RSA Benchmark for evaluating Audio-LLM robustness across six different tasks.

of *Qwen2-Audio-7B-Instruct* to reach **55.74%**, significantly worse than the **14.63** achieved with raw noisy audio. While modern methods like *DeepFilterNet* demonstrate a better ability to preserve speech information, they still fail to beat the original noisy baseline. In the $K = 4$ Classroom scenario, *DeepFilterNet* results in a WER of **32.62** for *Qwen2*, remaining higher than the raw performance of **15.70**. This suggests that even advanced denoising tools struggle to preserve the specific acoustic features that ALLMs rely on.

7 case study

To investigate the limitations of current models, we analyze specific failure instances across different modalities in Figure 3. In the audio processing task, the model exhibited severe hallucination by incorrectly identifying a male speaker as female and fabricating details about “sneezing” and health conditions, leading to a score of 0.0. Similarly, in the textual QA task regarding the spread of the Theses, the model failed to address the specific constraint (duration) and instead provided a verbose historical context, missing the core answer (“two weeks”) entirely. These examples highlight critical challenges in cross-modal grounding and precise instruction following, where models struggle to distinguish between generating related knowledge and strictly adhering to user queries.

8 Conclusion

This study empirically confirms that current ALLMs lack the intrinsic robustness required for intricate real-world acoustic ecologies. We observe a severe functional collapse in cognitive reasoning tasks under complex acoustic interference, in stark contrast to their relatively stable perceptual capabilities. Furthermore, our experiments reveal that external speech enhancement strategies often exacerbate performance errors. To this end, investigating noise-aware instruction tuning or adversarial training paradigms is essential for cultivating stability against environmental complexity.

570 Limitations

571 While this work provides a comprehensive diag-
572 nosis of ALLM vulnerabilities, our investigation
573 into improving robustness is limited to inference-
574 time mitigation via external speech enhancement.
575 Our results indicate that such "plug-and-play" pre-
576 processing often fails due to the model's sensitivity
577 to denoising artifacts. Consequently, we did not ex-
578 plore training-time interventions. Future research
579 should move beyond external patching and investi-
580 gate noise-aware instruction tuning or adversarial
581 training paradigms to cultivate intrinsic robustness
582 within the models themselves.

583 References

584 Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkin-
585 son, Hany Awadalla, Nguyen Bach, Jianmin Bao,
586 Alon Benhaim, Martin Cai, Vishrav Chaudhary, Cong-
587 cong Chen, and 1 others. 2025. Phi-4-mini technical
588 report: Compact yet powerful multimodal lan-
589 guage models via mixture-of-loras. *arXiv preprint*
590 *arXiv:2503.01743*.

591 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama
592 Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
593 Diogo Almeida, Janko Altenschmidt, Sam Altman,
594 Shyamal Anadkat, and 1 others. 2023. Gpt-4 techni-
595 cal report. *arXiv preprint arXiv:2303.08774*.

596 Orevaoghene Ahia, Martijn Bartelds, Kabir Ahuja,
597 Hila Gonen, Valentin Hofmann, Siddhant Arora,
598 Shuyue Stella Li, Vishal Puttagunta, Mofetoluwa
599 Adeyemi, Charishma Buchireddy, and 1 others. 2025.
600 Blab: Brutally long audio bench. *arXiv preprint*
601 *arXiv:2505.03054*.

602 Harsh Ahlawat, Naveen Aggarwal, and Deepti Gupta.
603 2025. Automatic speech recognition: A survey of
604 deep learning techniques and approaches. *Internation-*
605 *al Journal of Cognitive Computing in Engineer-*
606 *ing*.

607 Joseph Ozgis Akomodi, Maddisyn Neita, Samia
608 Ahmed, Aiyana Parnell, Aaliyah Taylor, Aaliya Tay-
609 lor, and Neita Maddisyn. 2025. Statistical data analy-
610 sis of signal-to-noise ratios in mri and ct scans. *Inter-*
611 *national Journal of Recent Innovations in Academic*
612 *Research*, 9(3):240–250.

613 MA Ali and PM Shemi. 2015. An improved method of
614 audio denoising based on wavelet transform. In *2015*
615 *international conference on Power, Instrumentation,*
616 *Control and Computing (PICC)*, pages 1–6. IEEE.

617 Hanin Atwany, Abdul Waheed, Rita Singh, Monojit
618 Choudhury, and Bhiksha Raj. 2025. Lost in tran-
619 scription, found in distribution shift: Demystifying
620 hallucination in speech foundation models. *arXiv*
621 *preprint arXiv:2502.12414*.

Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe
622 Kazemzadeh, Emily Mower, Samuel Kim, Jean-
623 nette N Chang, Sungbok Lee, and Shrikanth S
624 Narayanan. 2008. Iemocap: Interactive emotional
625 dyadic motion capture database. *Language resources*
626 *and evaluation*, 42(4):335–359. 627

Guangke Chen, Fu Song, Zhe Zhao, Xiaojun Jia,
628 Yang Liu, Yanchen Qiao, and Weizhe Zhang. 2025.
629 Audiojailbreak: Jailbreak attacks against end-to-
630 end large audio-language models. *arXiv preprint*
631 *arXiv:2505.14103*. 632

Sujal Chondhekar, Vasanth Murukuri, Rushabh Vasani,
633 Sanika Goyal, Rajshree Badami, Anushree Rana,
634 Sanjana SN, Karthik Pandia, Sulabh Katiyar, Neha
635 Jagadeesh, and 1 others. 2025. When de-noising
636 hurts: A systematic study of speech enhancement ef-
637 fects on modern medical asr systems. *arXiv preprint*
638 *arXiv:2512.17562*. 639

Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei,
640 Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng
641 He, Junyang Lin, and 1 others. 2024. Qwen2-audio
642 technical report. *arXiv preprint arXiv:2407.10759*. 643

Mohsen Fatehifar, Josef Schlittenlacher, Ibrahim Almu-
644 farrij, David Wong, Tim Cootes, and Kevin J Munro.
645 2025. Applications of automatic speech recognition
646 and text-to-speech technologies for hearing assess-
647 ment: a scoping review. *International Journal of*
648 *Audiology*, 64(6):537–548. 649

Arushi Goel, Sreyan Ghosh, Jaehyeon Kim, Sonal Ku-
650 mar, Zhifeng Kong, Sang-gil Lee, Chao-Han Huck
651 Yang, Ramani Duraiswami, Dinesh Manocha, Rafael
652 Valle, and 1 others. 2025. Audio flamingo 3: Advanc-
653 ing audio intelligence with fully open large audio
654 language models. *arXiv preprint arXiv:2507.08128*. 655

Xinlu He and Jacob Whitehill. 2025. Survey of end-to-
656 end multi-speaker automatic speech recognition for
657 monaural audio. *arXiv preprint arXiv:2505.10975*. 658

Yingxu He, Zhuohan Liu, Shuo Sun, Bin Wang, Wenyu
659 Zhang, Xunlong Zou, Nancy F Chen, and Ai Ti Aw.
660 2024. Meralion-audiollm: Bridging audio and lan-
661 guage with large language models. *arXiv preprint*
662 *arXiv:2412.09818*. 663

Gongping Huang, Jesper R Jensen, Jingdong Chen,
664 Jacob Benesty, Mads G Christensen, Akihiko
665 Sugiyama, Gary Elko, and Tomas Gaensler. 2025.
666 Advances in microphone array processing and multi-
667 channel speech enhancement. In *ICASSP 2025-2025*
668 *IEEE International Conference on Acoustics, Speech*
669 *and Signal Processing (ICASSP)*, pages 1–5. IEEE. 670

Rongjie Huang, Mingze Li, Dongchao Yang, Jia-
671 tong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu,
672 Zhiqing Hong, Jiawei Huang, Jinglin Liu, and 1 oth-
673 ers. 2024. Audiogpt: Understanding and generating
674 speech, music, sound, and talking head. In *Proceed-*
675 *ings of the AAAI Conference on Artificial Intelligence*,
676 volume 38, pages 23802–23804. 677

678	Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. <i>arXiv preprint arXiv:2410.21276</i> .	Zifan Peng, Yule Liu, Zhen Sun, Mingchen Li, Zeren Luo, Jingyi Zheng, Wenhan Dong, Xinlei He, Xuechao Wang, Yingjie Xue, and 1 others. 2025. Jalmbench: Benchmarking jailbreak vulnerabilities in audio language models. <i>arXiv preprint arXiv:2505.17568</i> .	735 736 737 738 739 740
683	Chaitanya Jannu and Sunny Dayal Vanambathina. 2025. An overview of speech enhancement based on deep learning techniques. <i>International Journal of Image and Graphics</i> , 25(01):2550001.	Karol J Piczak. 2015. Esc: Dataset for environmental sound classification. In <i>Proceedings of the 23rd ACM international conference on Multimedia</i> , pages 1015–1018.	741 742 743 744
687	Sonal Kumar, Šimon Sedláček, Vaibhavi Lokegaonkar, Fernando López, Wenyi Yu, Nishit Anand, Hyeonggon Ryu, Lichang Chen, Maxim Plička, Miroslav Hlaváček, and 1 others. 2025. Mmau-pro: A challenging and comprehensive benchmark for holistic evaluation of audio general intelligence. <i>arXiv preprint arXiv:2508.13992</i> .	Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. Meld: A multimodal multi-party dataset for emotion recognition in conversations. In <i>Proceedings of the 57th annual meeting of the association for computational linguistics</i> , pages 527–536.	745 746 747 748 749 750
694	Chen-An Li, Tzu-Han Lin, and Hung-yi Lee. 2025a. When silence matters: The impact of irrelevant audio on text reasoning in large audio-language models. <i>arXiv preprint arXiv:2510.00626</i> .	Tim Sainburg, Marvin Thielk, and Timothy Q Gentner. 2020. Finding, visualizing, and quantifying latent structure across diverse animal vocal repertoires. <i>PLoS computational biology</i> , 16(10):e1008228.	751 752 753 754
698	Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhat-tacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, and 1 others. 2025b. From generation to judgment: Opportunities and challenges of llm-as-a-judge. In <i>Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing</i> , pages 2757–2791.	Mohammad Sarim, Saim Shakeel, Laeaba Javed, Mohammad Nadeem, and 1 others. 2025. Direct speech to speech translation: A review. <i>arXiv preprint arXiv:2503.04799</i> .	755 756 757 758
706	Chaoqun Liu, Mahani Aljunied, Guizhen Chen, Hou Pong Chan, Weiwen Xu, Yu Rong, and Wenxuan Zhang. 2025a. Seallms-audio: Large audio-language models for southeast asia. <i>arXiv preprint arXiv:2511.01670</i> .	Hendrik Schröter, Tobias Rosenkranz, Alberto N Escalante-B, and Andreas Maier. 2023. Deepfilternet: Perceptually motivated real-time speech enhancement. <i>arXiv preprint arXiv:2305.08227</i> .	759 760 761 762
711	Yao Liu, Faizahani binti Ab Rahman, and Farah binti Mohamad Zain. 2025b. A systematic literature review of research on automatic speech recognition in efl pronunciation. <i>Cogent Education</i> , 12(1):2466288.	Suwon Shon, Siddhant Arora, Chyi-Jiunn Lin, Ankita Pasad, Felix Wu, Roshan Sharma, Wei-Lun Wu, Hung-yi Lee, Karen Livescu, and Shinji Watanabe. 2023. Slue phase-2: A benchmark suite of diverse spoken language understanding tasks. In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 8906–8937.	763 764 765 766 767 768 769 770
715	Ziyang Ma, Yinghao Ma, Yanqiao Zhu, Chen Yang, Yi-Wen Chao, Ruiyang Xu, Wenxi Chen, Yuanzhe Chen, Zhuo Chen, Jian Cong, and 1 others. 2025. Mmar: A challenging benchmark for deep reasoning in speech, audio, music, and their mix. <i>arXiv preprint arXiv:2505.13032</i> .	Jinhui Song, Xingxing Shi, Jiu Hui Wu, Tengyue Zheng, and Zhiwei Song. 2025. Nonlinear multi-order coupled stochastic resonance modeling under extremely low signal-to-noise ratios. <i>Mechanical Systems and Signal Processing</i> , 224:112208.	771 772 773 774 775
721	Daniela Silvia Pace, Giulia Pedrazzi, Irene D’amario, Aurora Troccoli, Giancarlo Giacomini, Maria Silvia Labriola, Gianni Pavan, Daniele Ventura, Edoardo Casoli, Giandomenico Ardizzone, and 1 others. 2025. The acoustic ecology of coastal dolphins by assessing the structural variability of sounds and the influence of contextual factors. <i>Integrative Zoology</i> , 20(4):686–699.	Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2023. Salmonn: Towards generic hearing abilities for large language models. <i>arXiv preprint arXiv:2310.13289</i> .	776 777 778 779 780
729	Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In <i>2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)</i> , pages 5206–5210. IEEE.	Jean-Marc Valin. 2018. A hybrid dsp/deep learning approach to real-time full-band speech enhancement. In <i>2018 IEEE 20th international workshop on multimedia signal processing (MMSP)</i> , pages 1–5. IEEE.	781 782 783 784
730		Bin Wang, Xunlong Zou, Geyu Lin, Shuo Sun, Zhuohan Liu, Wenyu Zhang, Zhengyuan Liu, AiTi Aw, and Nancy Chen. 2025a. Audiobench: A universal benchmark for audio large language models. In <i>Proceedings of the 2025 Conference of the Nations of</i>	785 786 787 788 789

790		Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, and 1 others. 2024. Minicpm-v: A gpt-4v level mllm on your phone. <i>arXiv preprint arXiv:2408.01800</i> .	843
791			844
792			845
793			846
794			847
795	Wei Wang, Siyi Zhao, and Yanmin Qian. 2025b. Advancing non-intrusive suppression on enhancement distortion for noise robust asr. In <i>ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 1–5. IEEE.		848
796			849
797			850
798			851
799	Chengwei Wei, Bin Wang, Jung-jae Kim, and Nancy F Chen. 2025. Towards spoken mathematical reasoning: Benchmarking speech-based models over multi-faceted math problems. <i>arXiv preprint arXiv:2505.15000</i> .		852
800			853
801			854
802			855
803			856
804	Kendall Wrightson. 2000. An introduction to acoustic ecology. <i>Soundscape: The journal of acoustic ecology</i> , 1(1):10–13.		857
805			858
806			859
807	Boyong Wu, Chao Yan, Chen Hu, Cheng Yi, Chengli Feng, Fei Tian, Feiyu Shen, Gang Yu, Haoyang Zhang, Jingbei Li, and 1 others. 2025. Step-audio 2 technical report. <i>arXiv preprint arXiv:2507.16632</i> .		860
808			861
809			862
810			863
811	Zhifei Xie, Mingbao Lin, Zihang Liu, Pengcheng Wu, Shuicheng Yan, and Chunyan Miao. 2025. Audio-reasoner: Improving reasoning capability in large audio language models. <i>arXiv preprint arXiv:2503.02318</i> .		864
812			865
813			866
814			867
815			868
816	Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, and 1 others. 2025a. Qwen2. 5-omni technical report. <i>arXiv preprint arXiv:2503.20215</i> .		869
817			870
818			871
819			872
820	Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfu Zhu, Yuanjun Lv, Yongqi Wang, Dake Guo, He Wang, Linhan Ma, Pei Zhang, Xinyu Zhang, Hongkun Hao, Zishan Guo, and 19 others. 2025b. Qwen3-omni technical report. <i>arXiv preprint arXiv:2509.17765</i> .		873
821			874
822			875
823			876
824			877
825			878
826			879
827	Chih-Kai Yang, Neo S Ho, and Hung-yi Lee. 2025a. Towards holistic evaluation of large audio-language models: A comprehensive survey. <i>arXiv preprint arXiv:2505.15957</i> .		880
828			881
829			882
830			883
831	Hao Yang, Lizhen Qu, Ehsan Shareghi, and Gholamreza Haffari. 2025b. Audio is the achilles’ heel: Red teaming audio large multimodal models. In <i>Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 9292–9306.		884
832			885
833			886
834			887
835			888
836			889
837			890
838	Qian Yang, Jin Xu, Wenrui Liu, Yunfei Chu, Ziyue Jiang, Xiaohuan Zhou, Yichong Leng, Yuanjun Lv, Zhou Zhao, Chang Zhou, and 1 others. 2024. Air-bench: Benchmarking large audio-language models via generative comprehension. <i>arXiv preprint arXiv:2402.07729</i> .		891
839			892
840			893
841			894
842			895

A Details of Speech Enhancement Algorithms

To investigate potential mitigation strategies against acoustic robustness degradation, we employed four distinct speech enhancement baselines. These methods range from traditional signal processing to state-of-the-art deep learning architectures, allowing us to evaluate the impact of different denoising paradigms on ALLM perception.

Noisereduce. This algorithm represents a traditional baseline based on stationary spectral gating. It operates by computing the Short-Time Fourier Transform of the noisy signal $Y(t, f)$ and applying a frequency-domain mask $\mathcal{M}(t, f)$. Formally, the binary mask generation and the subsequent signal reconstruction are defined as:

$$\mathcal{M}(t, f) = \mathbb{1}(|Y(t, f)| > \mu_N(f) + \lambda\sigma_N(f)) \quad (5)$$

$$\hat{X}(t, f) = Y(t, f) \cdot \mathcal{M}(t, f) \quad (6)$$

where $\mathbb{1}(\cdot)$ is the indicator function, μ_N and σ_N denote the mean and standard deviation of the pre-calculated noise profile, and λ serves as a sensitivity threshold. While effective for constant background noise, this hard-thresholding approach is non-learnable and prone to introducing spectral artifacts.

RNNoise. Designed for real-time VoIP, RNNoise utilizes a hybrid architecture that combines classic signal processing with Recurrent Neural Networks. Instead of generating raw waveforms, the model predicts gain adjustments for distinct frequency bands. This method prioritizes human perceptual intelligibility, which may not strictly align with the feature extraction requirements of ALLMs.

Audio-Denoising. This method adopts a generative, fully convolutional approach based on a 1D U-Net architecture. Leveraging an encoder-decoder structure with skip connections, it performs end-to-end regression directly on the time-domain waveform. This represents a pure deep learning approach aimed at mapping noisy inputs to clean signals without relying on traditional DSP constraints.

DeepFilterNet. As a state-of-the-art low-complexity framework, DeepFilterNet leverages complex-valued deep filtering techniques. It models spectral envelopes using Deep Neural Networks within the Equivalent Rectangular Bandwidth domain. This advanced architecture allows for superior handling of non-stationary noise, serving as

a high-fidelity benchmark to test whether sophisticated reconstruction can preserve the semantic cues required by ALLMs.

B More specific details Example

In this section, we present additional examples and detailed experimental data.

919
920
921

922

923

924

Evaluation Case Study

Question: How long did it take for the Theses printing to spread throughout Germany?

Reference: *two weeks*

Model Prediction:

The Theses of Martin Luther, which were printed and distributed in Wittenberg on October 31, 1517, quickly spread throughout Germany. Within a few months, they had reached many parts of the Holy Roman Empire, including major cities like Nuremberg and Augsburg. The rapid dissemination was facilitated by the advent of printing technology, particularly the invention of movable type by Johannes Gutenberg around that time. This allowed for more copies to be produced and circulated much faster than through handwritten manuscripts alone.

The widespread distribution of these theses played a significant role in sparking the Protestant Reformation, as it brought about increased debate among scholars and clergy regarding issues such as papal authority, indulgences, and church practices.

Judge Response (Score: 1.0):

Explanation: The reference answer states that it took "two weeks" for the Theses to spread throughout Germany, while the model's answer provides a detailed explanation of the context and significance of the Theses' dissemination but does not mention the specific timeframe of two weeks. Therefore, the model's response diverges significantly from the reference in providing an accurate answer to the question asked. It does not address the duration explicitly, focusing instead on the impact and the means of dissemination. I think this results in a low alignment with the reference answer regarding the specific detail of the timeframe.

Evaluation Case Study

Question: Consider a connected graph G with seven vertices and ten edges. Determine the genus of the graph.

Reference: The genus g of a graph G with v vertices and e edges can be determined using Euler's formula for a connected graph embedded on a surface of genus g : $v - e + f = 2 - 2g$, where f is the number of faces.

For G with $v = 7$ and $e = 10$, we solve for f : $7 - 10 + f = 2 - 2g \implies f = 5 + 2g$.

Substitute back into Euler's formula: $7 - 10 + (5 + 2g) = 2 - 2g \implies 4g = 0 \implies g = 0$.

The genus of the graph G is 0 (planar).

Model Prediction:

A connected graph with 7 vertices and 10 edges is a tree. In graph theory, a tree is the simplest type of connected graph where there are no cycles. A tree with n vertices has exactly $n - 1$ edges, so in this case, it's a tree because $7 - 1 = 6$, which equals the number of edges (10).

Judge Response (Score: 0.0):

Explanation: The reference answer accurately applies Euler's formula to determine the genus of the graph, concluding that the genus is 0, indicating that the graph can be embedded on a plane or a sphere. In contrast, the model's answer incorrectly states that the graph is a tree, which is not possible given that a tree with 7 vertices would have only 6 edges, not 10. Thus, the model's answer provides incorrect information and does not address the genus calculation or the specific characteristics of the graph as outlined in the reference. Therefore, the model's answer is completely misaligned with the reference answer.

Evaluation Case Study

Question: From the audio, can you guess the speaker's gender (Male or Female)?

Reference: *Yes, I can analyze the audio and determine there is a male speaker.*

Model Prediction:

According to the content of this speaker, it seems that the gender of the speaker is female, and the speaker is in the state of sneeze. What's more, I guess that the speaker is a Youth. The health condition sounds healthy.

Judge Response (Score: 0.0):

Explanation: The reference answer states that the speaker is male, while the model's answer incorrectly identifies the speaker as female and also includes irrelevant details about sneezing and health condition. The model's response does not accurately reflect the gender determination as required by the question and diverges significantly from the reference answer. Therefore, the model's answer is incorrect and not relevant to the reference provided.

Evaluation Case Study (WER Metric)

Reference Text:

why fades the lotus of the water

Model Prediction:

the lotus of the water fades because it is a metaphor for how life is beauty can be fleeting

Metric Evaluation:

Metric: Word Error Rate (WER)

Score: 2.29 (228.57%)

Note: A WER > 1.0 indicates the prediction contains significantly more errors (insertions/substitutions) than the total length of the reference.

Evaluation Case Study (MathQA)

Question: what's 33.4 minus 32.56

Instruction: Perform the required arithmetic operation and provide the result.

Reference: 0.84

Model Prediction:

To find the result of 33.4 minus 32.56, follow these steps:

1. Subtract the numbers: - Line up the decimal points - Subtract each digit from right to left

Here's a step-by-step breakdown:

$$\begin{array}{r} 33.40 \\ - 32.56 \\ \hline \end{array}$$

2. Start subtracting from the rightmost digits: - The tenths place (0.4) is greater than the hundredths place (0.6), so no borrowing is needed. - $4 - 6 = -2$, but since we can't have a negative number in this position, borrow 1 from the ones place.

... (omitted specific incorrect borrowing logic for brevity) ...

So the final calculation results in:

$$0.83$$

Therefore, 33.4 minus 32.56 equals 0.83.

Judge Response (Score: 0.0):

Explanation: The user requested the result of the subtraction $33.4 - 32.56$. The correct mathematical result is 0.84. The model provides a detailed step-by-step derivation but makes calculation errors in the borrowing process, ultimately arriving at an incorrect answer of 0.83. Since the core task was a precise arithmetic operation and the result is wrong, the response is considered a failure.

Evaluation Case Study (MathQA)

Question: what's 33.4 minus 32.56

Instruction: Perform the required arithmetic operation and provide the result.

Reference: 0.84

Model Prediction:

To find the result of 33.4 minus 32.56, follow these steps:

1. Subtract the numbers: - Line up the decimal points - Subtract each digit from right to left

Here's a step-by-step breakdown:

$$\begin{array}{r} 33.40 \\ - 32.56 \\ \hline \end{array}$$

2. Start subtracting from the rightmost digits: - The tenths place (0.4) is greater than the hundredths place (0.6), so no borrowing is needed. - $4 - 6 = -2$, but since we can't have a negative number in this position, borrow 1 from the ones place.

... (omitted repetitive steps for brevity) ...

So the final calculation results in:

$$0.83$$

Therefore, 33.4 minus 32.56 equals 0.83.

Judge Response (Score: 0.0):

Explanation: The user requested the arithmetic operation $33.4 - 32.56$. The correct result is 0.84. The model attempted to perform the subtraction step-by-step but made calculation errors in the borrowing process (incorrectly calculating the final digit), leading to an incorrect result of 0.83. Since the answer is mathematically wrong, the response is rated as a failure.

«CASE_TITLE» fonttitle

Question: «QUESTION»

Reference (Judge Prompt):

[Reference Answer]
«REF_ANSWER»

[Model Answer]
«MODEL_ANSWER»

[Question]
«QUESTION»

[Task]
You are a strict evaluator. Rate the model's answer based on its alignment with the reference answer, focusing on accuracy and relevance to the reference. Be critical on details. If the model response is something like "cannot decide" / "unable to answer", rate as 0.

[Scoring Rubric]
Score 0: Refusal or non-committal (e.g., "cannot decide"), or no concrete result.
Score 0: Completely misaligned; incorrect or irrelevant compared to the reference.
Score 1: Minimal alignment; largely misunderstands or provides irrelevant details.
Score 2: On-topic but diverges significantly from the reference in accuracy or relevance.
Score 3: Generally aligned but missing key details or containing minor inaccuracies.
Score 4: Mostly accurate and relevant; closely follows the reference but could be clearer or more complete.
Score 5: Highly accurate and detailed; matches the reference answer essentially perfectly.

[Output Format]
Explanation: (Briefly compare the reference vs. the model answer and justify the score.)
Rating: (An integer in {0,1,2,3,4,5})

Model Prediction (filled to judge):

«MODEL_PREDICTION»

Judge Response (Score: «SCORE»):

«JUDGE_EXPLANATION»

Scene: Pasture										
Model	ASR					MR				
	K = 0	K = 1	K = 2	K = 3	K = 4	K = 0	K = 1	K = 2	K = 3	K = 4
Qwen2	3.45	4.67	8.08	14.48	24.80	66.00	58.00	39.00	20.00	10.00
SALMONN	10.49	16.16	28.38	73.41	161.12	18.00	8.00	5.00	3.00	0.00
SeaLLMs	5.52	19.05	15.78	36.55	65.47	62.00	52.00	31.00	22.00	9.00
Phi	1.67	2.65	5.35	12.80	25.97	3.00	5.00	1.00	1.00	0.00
MERaLION	2.34	5.32	12.74	22.49	38.13	74.00	67.00	48.00	28.00	15.00
Step2	3.90	5.27	7.81	14.10	25.94	75.00	66.00	48.00	30.00	14.00
MiniCPM	2.95	7.18	18.64	38.56	68.03	75.00	65.00	38.00	19.00	7.00
Qwen-Turbo	23.78	25.10	30.47	40.69	55.09	88.00	66.00	33.00	20.00	14.00
Qwen2.5-Omni	23.32	25.71	30.09	43.17	56.07	89.00	58.00	37.00	23.00	14.00
Qwen3-Omni	1.72	2.25	5.12	10.61	21.96	91.00	86.00	64.00	44.00	25.00
gpt-mini	50.01	60.46	72.85	87.32	100.04	93.00	74.00	50.00	29.00	15.00
Model	ER					SQA				
	K = 0	K = 1	K = 2	K = 3	K = 4	K = 0	K = 1	K = 2	K = 3	K = 4
Qwen2	51.53	34.21	30.96	29.69	28.43	79.85	78.63	76.03	75.25	73.38
SALMONN	40.54	27.62	25.63	25.13	25.13	79.90	75.69	70.39	68.92	65.34
SeaLLMs	47.81	22.64	21.46	20.46	20.27	78.58	76.86	75.05	70.59	65.44
Phi	49.92	23.52	23.10	21.53	22.49	85.74	84.36	83.97	82.79	81.08
MERaLION	52.61	56.21	53.33	50.46	48.74	80.69	81.27	81.23	73.68	78.92
Step2	56.82	37.70	35.75	34.94	32.80	81.37	80.54	77.01	75.69	72.94
MiniCPM	55.33	28.35	27.82	26.67	26.25	82.94	83.87	82.35	79.75	76.96
Qwen-Turbo	52.99	23.68	23.56	19.96	15.63	82.50	84.17	83.33	79.31	75.00
Qwen2.5-Omni	52.91	24.10	24.06	19.50	18.47	83.82	84.90	82.16	79.02	75.88
Qwen3-Omni	47.20	32.03	29.81	28.47	26.55	80.98	81.52	82.50	81.76	80.10
gpt-mini	30.61	5.90	5.17	1.88	92.00	86.62	85.83	86.37	86.32	85.98
Model	GR					SI				
	K = 0	K = 1	K = 2	K = 3	K = 4	K = 0	K = 1	K = 2	K = 3	K = 4
Qwen2	96.02	94.52	94.82	95.12	94.12	49.60	46.60	45.80	35.20	30.40
SALMONN	82.37	84.86	81.87	76.29	76.29	58.40	53.00	55.00	53.20	56.00
SeaLLMs	79.87	72.73	76.01	69.62	72.62	62.00	41.40	37.60	26.60	20.20
Phi	38.65	43.13	36.35	31.47	24.20	33.20	17.40	20.00	22.40	18.00
MERaLION	85.26	84.36	83.86	82.47	78.98	71.00	68.80	62.40	55.00	42.60
Step2	86.95	88.45	83.37	83.76	77.19	58.20	50.80	44.80	40.80	35.20
MiniCPM	93.43	90.04	90.24	87.75	84.66	72.40	71.80	66.20	58.60	41.00
Qwen-Turbo	91.63	93.92	92.13	92.43	91.14	78.20	69.00	66.80	56.00	43.00
Qwen2.5-Omni	91.53	92.03	92.83	91.43	90.34	76.60	71.20	67.80	53.80	45.20
Qwen3-Omni	95.92	96.31	95.52	95.22	93.43	82.60	69.80	69.20	62.20	46.60
gpt-mini	-	-	-	-	-	78.20	79.20	68.60	59.40	38.20
Scene: Extreme Weather										
Model	ASR					MR				
	K = 0	K = 1	K = 2	K = 3	K = 4	K = 0	K = 1	K = 2	K = 3	K = 4
Qwen2	3.45	6.90	14.06	24.64	38.99	66.00	51.00	29.00	15.00	3.00
SALMONN	10.49	21.85	67.26	1.70	3.47	18.00	3.00	0.00	0.00	0.00
SeaLLMs	5.52	16.84	44.96	64.75	1.43	62.00	41.00	17.00	8.00	5.00
Phi	1.67	6.27	16.62	34.08	42.85	3.00	3.00	2.00	1.00	1.00
MERaLION	2.34	8.79	20.99	38.33	56.16	74.00	58.00	39.00	15.00	7.00
Step2	3.90	8.00	14.72	25.68	38.88	75.00	55.00	35.00	14.00	4.00
MiniCPM	2.95	14.87	33.43	64.42	80.03	75.00	49.00	30.00	10.00	5.00
Qwen-Turbo	24.00	25.00	41.00	47.00	66.00	88.00	59.00	26.00	13.00	7.00
Qwen2.5	85.27	92.19	1.02	1.28	1.51	61.00	46.00	26.00	15.00	3.00
Qwen2.5-Omni	23.00	25.00	37.00	50.00	66.00	89.00	52.00	35.00	10.00	7.00
Qwen3-Omni	2.00	5.00	12.00	43.00	95.00	91.00	75.00	51.00	29.00	14.00
gpt-mini	50.00	57.00	71.00	85.00	102.00	93.00	67.00	43.00	24.00	10.00
Model	ER					SQA				
	K = 0	K = 1	K = 2	K = 3	K = 4	K = 0	K = 1	K = 2	K = 3	K = 4
Qwen2	51.53	36.70	35.98	33.91	32.91	79.85	78.24	75.59	71.47	67.55
SALMONN	40.54	31.05	29.81	29.81	29.73	79.90	72.70	70.34	68.04	64.71
SeaLLMs	47.81	23.18	22.30	22.18	21.53	78.58	78.04	70.69	64.22	59.95
Phi	49.92	23.95	23.22	24.41	22.99	85.74	85.15	81.72	80.54	77.60
MERaLION	52.61	56.51	56.44	54.60	54.60	80.69	80.20	79.80	76.62	74.90
Step2	56.82	39.89	40.27	39.54	39.96	81.37	81.42	77.89	74.51	71.67
MiniCPM	55.33	29.27	30.46	30.38	30.84	82.94	82.65	81.72	76.52	76.86
Qwen-Turbo	52.99	23.87	24.56	17.82	12.49	82.50	82.30	80.78	76.32	74.46
Qwen2.5	18.07	16.45	17.87	17.87	16.03	75.25	72.75	70.15	69.07	65.44
Qwen2.5-Omni	52.91	23.45	25.56	17.66	14.75	83.82	83.04	81.27	80.10	75.59
Qwen3-Omni	47.20	33.98	33.98	35.40	36.40	80.98	81.37	78.33	81.47	78.87
gpt-mini	30.61	7.05	9.20	3.83	1.95	86.62	86.08	85.39	84.85	83.48
Model	GR					SI				
	K = 0	K = 1	K = 2	K = 3	K = 4	K = 0	K = 1	K = 2	K = 3	K = 4
Qwen2	96.02	95.02	92.83	92.93	92.93	49.60	45.40	31.20	29.00	13.40
SALMONN	82.37	83.27	75.40	69.22	64.04	58.40	53.00	51.80	54.20	56.80
SeaLLMs	classroom	69.34	73.99	69.42	66.73	62.00	46.40	34.40	23.40	8.60
Phi	38.65	40.14	39.74	30.08	22.21	33.20	23.60	19.40	20.60	13.00
MERaLION	85.26	83.86	83.07	84.16	79.78	71.00	64.60	56.20	45.60	27.00
Step2	86.95	89.34	85.26	83.76	79.38	58.20	49.60	40.80	37.80	23.40
MiniCPM	93.43	92.53	92.53	90.84	90.24	72.40	65.00	59.00	39.00	27.40
Qwen-Turbo	91.63	94.72	92.93	92.83	92.13	78.20	73.00	62.80	49.80	36.60
Qwen2.5	43.73	70.78	62.97	58.15	47.17	51.40	47.40	43.40	35.60	24.60
Qwen2.5-Omni	91.53	93.82	95.02	92.33	91.93	76.60	72.60	59.60	48.40	38.20

Continued on next page

<i>Qwen3-Omni</i>	95.92	95.42	95.62	95.22	94.22	82.60	75.60	62.60	57.40	41.80
<i>gpt-mini</i>	-	-	-	-	-	78.20	73.80	63.00	51.80	31.20
Scene: Classroom										
Model	ASR					MR				
	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4
<i>Qwen2</i>	3.45	4.24	6.47	9.95	14.63	66.00	62.00	46.00	33.00	26.00
<i>SALMONN</i>	10.49	11.74	16.85	27.14	35.53	18.00	11.00	6.00	4.00	3.00
<i>SeaLLMs</i>	5.52	7.42	15.08	27.63	52.31	62.00	57.00	47.00	35.00	27.00
<i>Phi</i>	1.67	2.32	3.67	8.81	12.67	<u>3.00</u>	<u>5.00</u>	<u>5.00</u>	<u>3.00</u>	<u>3.00</u>
<i>MERaLION</i>	2.34	3.46	6.10	10.34	15.70	74.00	71.00	69.00	52.00	41.00
<i>Step2</i>	3.90	5.67	7.56	8.66	11.68	75.00	70.00	61.00	50.00	38.00
<i>MiniCPM</i>	2.95	6.25	10.36	18.39	29.66	75.00	73.00	56.00	48.00	38.00
<i>Qwen-Turbo</i>	23.78	23.92	25.91	26.32	28.95	88.00	72.00	59.00	47.00	36.00
<i>Qwen2.5-Omni</i>	23.32	24.16	24.32	26.06	28.83	89.00	69.00	55.00	46.00	36.00
<i>Qwen3-Omni</i>	1.72	2.51	3.44	4.77	7.67	91.00	82.00	80.00	72.00	59.00
<i>gpt-mini</i>	<u>50.01</u>	<u>55.12</u>	<u>60.48</u>	<u>66.27</u>	<u>75.78</u>	93.00	83.00	67.00	51.00	43.00
Model	ER					SQA				
	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4
<i>Qwen2</i>	51.53	37.47	36.28	34.98	33.72	79.85	78.68	78.28	75.54	<u>72.45</u>
<i>SALMONN</i>	40.54	32.08	31.00	29.77	29.08	79.90	78.38	<u>75.20</u>	72.94	73.48
<i>SeaLLMs</i>	47.81	23.18	21.95	22.07	21.00	<u>78.58</u>	<u>77.21</u>	<u>77.75</u>	75.88	73.77
<i>Phi</i>	49.92	23.91	24.18	23.37	22.53	85.74	84.56	83.14	83.68	84.17
<i>MERaLION</i>	52.61	56.59	55.90	55.36	54.67	80.69	81.76	79.75	81.76	79.61
<i>Step2</i>	56.82	40.61	38.70	38.89	37.78	81.37	81.76	79.56	77.21	76.91
<i>MiniCPM</i>	55.33	30.27	28.97	27.82	27.24	82.94	83.19	82.60	83.58	81.32
<i>Qwen-Turbo</i>	52.99	26.21	26.63	22.99	21.34	82.50	84.46	83.82	83.97	82.25
<i>Qwen2.5-Omni</i>	52.91	26.93	27.43	23.37	22.18	83.82	84.71	83.63	83.68	81.27
<i>Qwen3-Omni</i>	47.20	33.49	32.26	32.80	32.80	80.98	80.93	81.72	82.30	79.71
<i>gpt-mini</i>	30.61	8.51	8.74	4.71	3.26	86.62	86.27	85.29	86.13	85.69
Model	GR					SI				
	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4
<i>Qwen2</i>	96.02	95.22	95.62	95.52	94.62	49.60	47.20	45.60	35.00	29.60
<i>SALMONN</i>	82.37	78.49	76.20	72.81	64.54	58.40	57.40	54.00	54.80	56.20
<i>SeaLLMs</i>	classroom	77.17	75.81	77.06	76.47	62.00	49.20	47.20	40.20	36.00
<i>Phi</i>	38.65	39.64	38.84	33.96	29.88	<u>33.20</u>	<u>21.40</u>	<u>20.00</u>	<u>20.40</u>	<u>19.60</u>
<i>MERaLION</i>	85.26	84.06	84.96	84.06	82.17	71.00	69.60	67.20	65.80	65.00
<i>Step2</i>	86.95	89.74	89.24	89.74	89.84	58.20	53.40	46.60	40.00	34.80
<i>MiniCPM</i>	93.43	92.93	90.14	89.14	85.86	72.40	72.00	67.80	67.80	63.00
<i>Qwen-Turbo</i>	91.63	94.02	94.02	94.32	94.22	78.20	73.60	76.00	69.40	64.80
<i>Qwen2.5-Omni</i>	91.53	94.52	94.62	95.22	93.13	76.60	71.80	70.80	72.80	66.80
<i>Qwen3-Omni</i>	95.92	96.31	96.12	95.92	94.12	82.60	74.20	72.80	77.00	73.40
<i>gpt-mini</i>	-	-	-	-	-	78.20	79.00	80.00	75.00	61.00
Scene: Outdoors										
Model	ASR					MR				
	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4
<i>Qwen2</i>	3.45	8.49	19.67	35.97	54.75	66.00	43.00	23.00	6.00	4.00
<i>SALMONN</i>	10.49	24.47	124.79	317.33	509.11	18.00	5.00	2.00	0.00	0.00
<i>SeaLLMs</i>	5.52	25.49	51.13	125.50	279.27	62.00	29.00	10.00	1.00	0.00
<i>Phi</i>	1.67	7.07	19.54	42.89	81.12	<u>3.00</u>	<u>1.00</u>	<u>2.00</u>	2.00	1.00
<i>MERaLION</i>	2.34	11.63	30.89	55.35	76.04	74.00	46.00	27.00	7.00	3.00
<i>Step2</i>	3.90	7.59	20.47	34.49	66.67	75.00	47.00	26.00	12.00	6.00
<i>MiniCPM</i>	2.95	21.08	57.34	89.09	121.17	75.00	42.00	18.00	7.00	1.00
<i>Qwen-Turbo</i>	23.78	27.95	42.30	61.54	96.42	88.00	48.00	18.00	7.00	4.00
<i>Qwen2.5-Omni</i>	23.32	28.89	45.18	61.57	93.27	89.00	39.00	16.00	4.00	4.00
<i>Qwen3-Omni</i>	1.72	5.70	48.41	259.56	<u>557.20</u>	91.00	63.00	40.00	18.00	5.00
<i>gpt-mini</i>	<u>50.01</u>	<u>64.69</u>	<u>86.97</u>	<u>107.27</u>	<u>118.39</u>	93.00	49.00	16.00	6.00	3.00
Model	ER					SQA				
	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4
<i>Qwen2</i>	51.53	35.33	33.68	32.84	30.84	79.85	77.21	73.97	67.21	62.25
<i>SALMONN</i>	40.54	30.87	30.46	30.46	30.23	79.90	73.14	67.84	63.82	62.75
<i>SeaLLMs</i>	47.81	22.45	21.00	21.23	20.23	78.58	73.92	65.93	58.38	55.74
<i>Phi</i>	49.92	24.14	23.72	24.10	23.95	85.74	86.42	80.69	80.05	73.28
<i>MERaLION</i>	52.61	56.59	55.82	56.51	55.06	80.69	80.29	77.60	73.68	71.08
<i>Step2</i>	<u>56.82</u>	39.27	38.58	38.39	39.69	81.37	79.31	76.47	69.07	61.27
<i>MiniCPM</i>	55.33	30.04	29.35	29.85	29.43	82.94	82.45	80.00	75.78	71.37
<i>Qwen-Turbo</i>	52.99	22.91	25.02	14.10	10.57	82.50	82.65	79.71	70.20	66.62
<i>Qwen2.5-Omni</i>	52.91	23.18	24.90	13.60	12.34	83.82	81.37	78.97	73.53	66.96
<i>Qwen3-Omni</i>	47.20	34.10	33.56	33.41	34.10	80.98	80.83	81.03	75.74	73.53
<i>gpt-mini</i>	30.61	<u>5.67</u>	8.93	1.07	<u>0.23</u>	86.62	86.18	86.18	83.38	79.17
Model	GR					SI				
	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4	<i>K</i> = 0	<i>K</i> = 1	<i>K</i> = 2	<i>K</i> = 3	<i>K</i> = 4
<i>Qwen2</i>	96.02	93.62	90.04	87.65	81.77	49.60	43.20	32.60	19.80	6.60
<i>SALMONN</i>	82.37	82.70	71.40	65.80	59.70	58.40	51.60	55.20	57.40	54.60
<i>SeaLLMs</i>	79.87	71.00	74.20	73.20	75.70	62.00	39.20	21.60	10.00	3.60
<i>Phi</i>	38.65	38.65	32.07	27.69	21.41	<u>33.20</u>	<u>26.40</u>	<u>18.40</u>	<u>17.60</u>	<u>11.80</u>
<i>MERaLION</i>	85.26	83.00	84.10	81.80	76.50	71.00	67.40	55.20	37.40	17.20
<i>Step2</i>	86.95	86.00	81.70	83.80	77.20	58.20	51.60	36.40	23.80	9.80
<i>MiniCPM</i>	93.43	91.30	90.10	85.20	81.90	72.40	59.80	46.80	30.20	12.00
<i>Qwen-Turbo</i>	91.63	91.53	89.84	88.55	86.35	78.20	70.40	53.80	36.20	21.00
<i>Qwen2.5-Omni</i>	91.53	91.04	90.14	89.54	87.05	76.60	68.60	56.60	34.20	20.00
<i>Qwen3-Omni</i>	95.92	92.53	91.33	88.35	88.94	82.60	70.20	58.40	41.40	17.80

Continued on next page

<i>gpt-mini</i>	-	-	-	-	-		78.20	76.00	52.00	29.60	6.40
-----------------	---	---	---	---	---	--	-------	--------------	-------	-------	------

Table 4: **Graph3 comprehensive results across four dimensions.** Blocks are scenes; rows are models; columns are noise-source count $K = 0 \dots 4$ for each task. Values are percentages (numbers < 1 are multiplied by 100). ASR reports WER (lower is better); others are higher-is-better. Best (worst) within each scene+task column is in **bold** (underlined).

Scene: Pasture									
ASR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	MR	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	4.67	8.08	14.48	24.80	Noise	58.00	39.00	20.00	10.00
NoiseReduce	12.60	28.39	51.31	73.63	NoiseReduce	<u>32.00</u>	17.00	<u>4.00</u>	1.00
AudioDenoise	7.28	17.21	34.17	56.87	AudioDenoise	45.00	30.00	11.00	3.00
PyRNNoise	<u>15.79</u>	34.59	<u>60.09</u>	81.43	PyRNNoise	35.00	9.00	5.00	<u>0.00</u>
DeepFilterNet	9.62	24.28	41.81	<u>60.68</u>	DeepFilterNet	50.00	25.00	7.00	3.00
ER	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SQA	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	34.21	30.96	29.69	28.43	Noise	78.63	76.03	75.25	73.38
NoiseReduce	31.23	<u>27.28</u>	29.43	22.95	NoiseReduce	76.18	73.33	68.43	61.37
AudioDenoise	<u>30.96</u>	28.66	28.24	<u>20.31</u>	AudioDenoise	78.09	74.71	69.12	63.77
PyRNNoise	31.88	30.08	<u>27.36</u>	29.43	PyRNNoise	<u>74.71</u>	<u>70.05</u>	<u>63.38</u>	<u>57.89</u>
DeepFilterNet	33.49	31.23	<u>28.66</u>	32.18	DeepFilterNet	79.36	75.74	<u>69.36</u>	64.46
GR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SI	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	94.52	94.82	95.12	94.12	Noise	46.60	45.80	35.20	30.40
NoiseReduce	88.94	87.35	81.97	81.18	NoiseReduce	<u>37.60</u>	<u>22.80</u>	16.80	4.60
AudioDenoise	95.72	93.43	91.33	88.94	AudioDenoise	42.80	34.40	23.40	15.20
PyRNNoise	91.83	88.25	82.67	<u>74.70</u>	PyRNNoise	38.60	26.40	<u>12.40</u>	<u>2.60</u>
DeepFilterNet	92.03	90.94	92.03	89.64	DeepFilterNet	41.80	31.40	20.00	12.60
Scene: Classroom									
ASR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	MR	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	4.24	6.47	9.95	14.63	Noise	62.00	46.00	33.00	26.00
NoiseReduce	<u>12.90</u>	<u>24.16</u>	<u>38.56</u>	<u>55.74</u>	NoiseReduce	<u>29.00</u>	<u>18.00</u>	<u>9.00</u>	<u>4.00</u>
AudioDenoise	5.62	10.88	18.67	30.61	AudioDenoise	55.00	37.00	27.00	14.00
PyRNNoise	7.73	15.08	24.62	36.48	PyRNNoise	46.00	33.00	21.00	16.00
DeepFilterNet	7.71	14.03	22.51	32.62	DeepFilterNet	44.00	32.00	21.00	18.00
Xhline1.1pt ER	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SQA	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	37.47	36.28	34.98	33.72	Noise	78.68	78.28	75.54	72.45
NoiseReduce	<u>32.11</u>	<u>29.16</u>	35.21	25.79	NoiseReduce	<u>76.13</u>	<u>73.48</u>	<u>70.20</u>	<u>63.77</u>
AudioDenoise	36.28	34.14	<u>34.21</u>	<u>23.64</u>	AudioDenoise	78.87	75.20	70.74	65.20
PyRNNoise	32.80	29.35	35.75	31.11	PyRNNoise	79.71	76.08	74.80	70.49
DeepFilterNet	34.56	32.34	34.83	32.64	DeepFilterNet	81.32	77.01	73.53	72.70
GR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SI	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	95.22	95.62	95.52	94.62	Noise	47.20	45.60	35.00	29.60
NoiseReduce	88.55	84.56	83.47	<u>79.68</u>	NoiseReduce	<u>33.20</u>	<u>29.20</u>	<u>21.60</u>	8.40
AudioDenoise	94.82	92.73	93.53	93.23	AudioDenoise	42.40	47.40	33.20	22.60
PyRNNoise	92.63	91.43	92.13	90.94	PyRNNoise	47.60	40.20	35.80	25.40
DeepFilterNet	92.23	92.83	93.43	92.23	DeepFilterNet	43.00	41.60	35.60	27.80

Table 5: Denoising mitigation for Qwen2 across two acoustic scenarios. The table reports Qwen2 performance under increasing multi-source acoustic interference ($K = 1 \dots 4$) in Pasture and Classroom. Each task block compares the no-denoise baseline (Noise) with four denoising methods; ASR is WER (lower is better) and other tasks are higher-is-better. **Bold/underline are computed among all methods** (including the Noise baseline).

Scene: Pasture									
ASR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	MR	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	5.32	12.74	22.49	38.13	Noise	67.00	48.00	28.00	15.00
NoiseReduce	<u>13.34</u>	30.96	54.26	73.05	NoiseReduce	40.00	<u>20.00</u>	6.00	3.00
AudioDenoise	8.27	20.36	39.06	61.20	AudioDenoise	52.00	34.00	13.00	6.00
PyRNNoise	12.58	31.89	56.58	78.60	PyRNNoise	38.00	21.00	5.00	1.00
DeepFilterNet	7.64	19.29	37.01	56.09	DeepFilterNet	52.00	28.00	8.00	9.00
ER	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SQA	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	56.21	53.33	50.46	48.74	Noise	81.27	81.23	73.68	78.92
NoiseReduce	51.42	50.65	47.43	44.48	NoiseReduce	78.43	80.39	75.00	70.74
AudioDenoise	52.76	50.65	47.16	<u>38.47</u>	AudioDenoise	80.98	79.71	77.75	71.47
PyRNNoise	49.46	<u>47.74</u>	46.93	43.52	PyRNNoise	78.53	75.44	<u>71.08</u>	<u>68.53</u>
DeepFilterNet	<u>47.93</u>	49.31	<u>46.55</u>	47.32	DeepFilterNet	81.13	<u>76.86</u>	<u>77.16</u>	<u>71.91</u>
GR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SI	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	84.36	83.86	82.47	78.98	Noise	68.80	62.40	55.00	42.60
NoiseReduce	83.07	81.47	75.20	67.13	NoiseReduce	64.60	48.60	31.00	15.00
AudioDenoise	83.86	80.88	76.89	67.13	AudioDenoise	69.20	59.40	44.60	24.40
PyRNNoise	<u>79.68</u>	<u>76.29</u>	<u>71.61</u>	<u>65.54</u>	PyRNNoise	<u>63.40</u>	49.00	<u>25.20</u>	<u>6.00</u>
DeepFilterNet	79.28	77.59	76.69	76.00	DeepFilterNet	66.40	60.80	39.00	24.80
Scene: Classroom									
ASR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	MR	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	3.46	6.10	10.34	15.70	Noise	71.00	69.00	52.00	41.00
NoiseReduce	<u>12.70</u>	<u>25.08</u>	41.62	60.09	NoiseReduce	44.00	27.00	13.00	8.00
AudioDenoise	4.88	9.93	18.56	<u>30.15</u>	AudioDenoise	67.00	53.00	29.00	23.00
PyRNNoise	5.53	11.36	19.84	29.82	PyRNNoise	53.00	38.00	28.00	19.00
DeepFilterNet	5.77	10.50	16.98	25.00	DeepFilterNet	60.00	44.00	28.00	19.00
ER	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SQA	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	56.59	55.90	55.36	54.67	Noise	81.76	79.75	81.76	79.61
NoiseReduce	51.15	51.30	52.34	<u>45.36</u>	NoiseReduce	80.29	<u>79.41</u>	<u>78.53</u>	<u>74.51</u>
AudioDenoise	52.87	54.37	53.10	<u>45.33</u>	AudioDenoise	80.83	81.52	80.93	78.38
PyRNNoise	51.84	51.19	51.92	49.35	PyRNNoise	80.78	80.39	79.41	78.28
DeepFilterNet	46.93	50.96	52.68	47.85	DeepFilterNet	<u>80.05</u>	79.51	80.93	77.70
GR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SI	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	84.06	84.96	84.06	82.17	Noise	69.60	67.20	65.80	65.00
NoiseReduce	84.76	80.48	78.78	<u>74.40</u>	NoiseReduce	<u>60.80</u>	<u>53.60</u>	<u>38.20</u>	<u>22.60</u>
AudioDenoise	84.06	83.37	78.78	<u>74.70</u>	AudioDenoise	69.60	64.20	60.40	52.80
PyRNNoise	78.49	<u>79.08</u>	80.78	78.39	PyRNNoise	68.00	66.00	60.40	52.60
DeepFilterNet	79.38	80.08	<u>77.89</u>	78.59	DeepFilterNet	67.80	65.20	63.20	50.00

Table 6: Denoising mitigation for MERaLION across two acoustic scenarios. The table reports MERaLION performance under increasing multi-source acoustic interference ($K = 1 \dots 4$) in Pasture and Classroom. Each task block compares the no-denoise baseline (Noise) with four denoising methods; ASR is WER (lower is better) and other tasks are higher-is-better. **Bold/underline are computed among all methods** (including the Noise baseline).

Scene: Pasture									
ASR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	MR	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	5.27	7.81	14.10	25.94	Noise	66.00	48.00	30.00	14.00
NoiseReduce	10.57	25.63	69.87	183.00	NoiseReduce	40.00	22.00	6.00	2.00
AudioDenoise	6.97	13.74	34.34	55.88	AudioDenoise	54.00	40.00	14.00	6.00
PyRNNNoise	14.69	37.58	73.92	105.67	PyRNNNoise	42.00	15.00	7.00	1.00
DeepFilterNet	10.06	20.20	46.42	86.65	DeepFilterNet	52.00	31.00	10.00	6.00
ER	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SQA	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	37.70	35.75	34.94	32.80	Noise	80.54	77.01	75.69	72.94
NoiseReduce	<u>35.10</u>	<u>32.34</u>	34.32	<u>23.64</u>	NoiseReduce	79.90	75.34	<u>67.75</u>	<u>58.43</u>
AudioDenoise	36.97	34.44	<u>31.65</u>	26.74	AudioDenoise	79.61	77.01	70.54	65.74
PyRNNNoise	37.89	35.63	33.60	32.45	PyRNNNoise	<u>79.12</u>	<u>74.71</u>	69.85	61.42
DeepFilterNet	38.66	37.13	32.45	39.04	DeepFilterNet	81.62	80.34	75.78	68.09
GR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SI	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	88.45	83.37	83.76	77.19	Noise	50.80	44.80	40.80	35.20
NoiseReduce	88.45	85.96	84.46	79.28	NoiseReduce	46.80	37.00	17.20	6.40
AudioDenoise	<u>86.25</u>	86.35	83.96	78.39	AudioDenoise	48.40	39.40	29.00	12.80
PyRNNNoise	88.94	84.96	<u>78.69</u>	<u>73.51</u>	PyRNNNoise	48.40	<u>25.60</u>	<u>13.40</u>	<u>2.80</u>
DeepFilterNet	86.95	<u>83.37</u>	80.58	77.89	DeepFilterNet	55.40	42.20	24.60	12.00
Scene: Classroom									
ASR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	MR	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	5.67	7.56	8.66	11.68	Noise	70.00	61.00	50.00	38.00
NoiseReduce	<u>12.19</u>	<u>21.02</u>	41.22	86.09	NoiseReduce	33.00	24.00	11.00	4.00
AudioDenoise	5.66	9.58	17.09	35.65	AudioDenoise	53.00	44.00	29.00	18.00
PyRNNNoise	8.21	10.69	20.78	30.57	PyRNNNoise	55.00	39.00	32.00	19.00
DeepFilterNet	7.03	13.07	20.20	28.51	DeepFilterNet	60.00	42.00	32.00	17.00
ER	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SQA	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	40.61	38.70	38.89	37.78	Noise	81.76	79.56	77.21	76.91
NoiseReduce	38.66	<u>37.32</u>	<u>35.79</u>	33.87	NoiseReduce	78.09	76.52	<u>69.80</u>	<u>67.79</u>
AudioDenoise	39.43	<u>37.32</u>	36.70	<u>26.74</u>	AudioDenoise	80.83	77.06	75.69	68.63
PyRNNNoise	<u>38.01</u>	<u>37.32</u>	36.09	34.60	PyRNNNoise	79.02	80.34	78.63	77.35
DeepFilterNet	38.54	38.54	37.09	37.89	DeepFilterNet	80.59	81.23	81.27	77.16
GR	$K = 1$	$K = 2$	$K = 3$	$K = 4$	SI	$K = 1$	$K = 2$	$K = 3$	$K = 4$
Noise	89.74	89.24	89.74	89.84	Noise	53.40	46.60	40.00	34.80
NoiseReduce	89.34	88.75	87.25	<u>83.86</u>	NoiseReduce	<u>46.40</u>	<u>38.60</u>	<u>22.20</u>	<u>11.40</u>
AudioDenoise	89.44	<u>87.85</u>	86.85	87.85	AudioDenoise	54.00	46.60	39.20	29.00
PyRNNNoise	91.14	88.05	85.86	86.85	PyRNNNoise	54.60	46.00	43.60	33.00
DeepFilterNet	<u>87.65</u>	89.34	<u>85.06</u>	85.96	DeepFilterNet	62.60	53.60	46.20	38.20

Table 7: Impact of denoising methods on Step2 under multi-source acoustic interference. The table compares six tasks (ASR, ER, GR, MR, SQA, SI) in two scenarios (Pasture vs. Classroom) as the number of interfering noise sources increases ($K = 1 \dots 4$). Each block reports the no-denoise baseline (Noise) and four denoising methods. For ASR we report WER (lower is better); other tasks use higher-is-better scores. **Bold/underline are computed among all methods** (including the Noise baseline).