

AU-Harness: An Open-Source Toolkit for Efficient and Unified Evaluation of Audio-LLMs

Hoang Nguyen[§], Sidharth Surapaneni^{†*}, Akshay Kalkunte[§], Jash Mehta[§], Aman Tiwari[§],
Oluwanifemi Bamgbose[§], Khyati Mahajan[§], Jash Shah[§], Shruthan Radhakrishna[§],
Sathwik Tejaswi Madhusudhan[§], Vikas Yadav[§], Sai Rajeswar[§]

[§] ServiceNow

[†] University of Texas at Austin

Correspondence: huhuhoang.nguyen@servicenow.com

Abstract

Large Audio Language Models (LALMs) are rapidly advancing, but evaluating them remains challenging due to inefficient and non-standardized toolkits that limit fair comparison and systematic assessment. Existing evaluation frameworks exhibit three critical limitations: (1) slow and inefficient processing pipeline that bottlenecks large-scale studies, (2) inadequate multi-turn dialogue support, leaving fundamental questions about cross-turn context integration and performance dynamics over extended conversations in LALMs unanswered; and (3) the absence of unified and scalable evaluation framework capable of keeping pace with the rapid growth of both LALMs and audio benchmarks. To address these issues, we introduce **AU-Harness**, an efficient and comprehensive evaluation framework for LALMs. Our empirical studies demonstrate that AU-Harness achieves 77% reduction in evaluation processing time over existing evaluation toolkits, enabling large-scale LALM evaluations previously considered impractical. We provide standardized evaluation protocols and flexible configurations for fair model comparison across diverse scenarios. AU-Harness unlocks a range of in-depth analyses difficult to conduct without a unified foundation, including multi-turn dialogue dynamics, enabling the study of true audio reasoning capabilities in existing LALMs. AU-Harness provides both practical evaluation tools and insights into model limitations, advancing systematic LALM development.¹

1 Introduction

The emergence of Large Audio Language Models (LALMs) has opened new frontiers, extending capabilities beyond textual inputs to speech, sounds, and multimodal inputs (Tang et al., 2023; Cui et al.,

2024). This progress has accelerated the development of frontier LALMs and audio-focused benchmarks. Recent multimodal LALMs like Gemini 2.5 (Comanici et al., 2025), Qwen2.5-Omni (Xu et al., 2025a) have demonstrated substantial audio understanding capabilities well beyond traditional Automatic Speech Recognition (ASR) tasks. However, despite these advances, audio evaluation toolkits² have comparatively received little attention. Thus, there is a need for efficient, customizable, and consistent evaluation framework for fair model comparisons which can evolve as audio tasks and model complexities grow.

Existing efforts including AudioBench (Wang et al., 2025a), Kimi-Eval (Ding et al., 2025), VoiceBench (Chen et al., 2024) and LMMS-Eval (Zhang et al., 2025b) have provided extensive task coverage from ASR to spoken question answering and scene understanding. However, prevailing toolkits still face three persistent limitations. First, **throughput**: many pipelines under-utilize batching and parallelism, creating bottlenecks that preclude large-scale, systematic comparisons. Second, **reproducibility**: ad-hoc prompting and non-standardized evaluation settings lead to incomparable performance across setups. Third, **task scope**: evaluations remain largely restricted to static single-turn interactions, failing to account for LALM assessment over extended interactions in multi-turn conversational settings which are more frequent in real world conversations.

Most current evaluation frameworks depend on simplistic yet inefficient input processing pipelines that struggle to scale with the increasing volume and complexity of audio benchmarks and LALMs. These limitations not only constrain the throughput of large-scale evaluations, but also hinder fair and reproducible comparisons across models of different sizes and architectures. As the field pro-

^{*}Work done during internship at ServiceNow

¹<https://github.com/ServiceNow/AU-Harness>

²We utilize the term evaluation **toolkit**, **kit**, and **framework** interchangeably across the manuscript.

EvalKit	HF Support	vLLM Support	Multi-task Parallel	Multi-turn	Customizable	RTF ↓	SPS ↑
AudioBench (Wang et al., 2025a)	✓	✗	✗	✗	✗	101.29 ± 6.37	0.34 ± 0.02
Kimi-Eval (Ding et al., 2025)	✓	✗	✗	✗	✗	54.13 ± 0.07	0.63 ± 0.01
VoiceBench (Chen et al., 2024)	✓	✗	✗	✗	✗	129.67 ± 5.30	0.29 ± 0.03
LMMS-Eval (Zhang et al., 2025b)	✓	✓	✓	✗	✗	20.84 ± 0.62	1.63 ± 0.05
AHELM (Lee et al., 2025)	✓	✓	✗	✓	✗	57.86 ± 14.69	0.62 ± 0.18
AU-Harness	✓	✓	✓	✓	✓	10.29 ± 0.61	3.32 ± 0.15

Table 1: **Feature comparison of contemporary LALM evaluation toolkits.** We evaluate key technical capabilities across existing frameworks: vLLM integration for efficient batching, HuggingFace (HF) model support, multi-turn dialogue support for conversational scenarios, parallel processing support for multi-task concurrent evaluation, and configurable customizations for flexible evaluation design. Throughput is further reported on MELD-Emotion (Poria et al., 2019) via Samples Per Second (SPS↑) and Real-Time Factor (RTF↓) averaged across 3 runs to empirically validate efficiency claims beyond feature support alone. Refer to §3.1 for further details. **Our framework provides comprehensive support across most dimensions.**

gresses toward more diverse and challenging audio tasks, the shortcomings of current evaluation infrastructure may pose a critical bottleneck, ultimately hampering the potential progress of LALMs. Unlike previous frameworks, we introduce an efficient token request orchestration together with effective data sharding to scale evaluations across multiple nodes and hardware architectures, leading to improved efficiency for audio benchmark evaluations.

Beyond computational efficiency, existing toolkits exhibit critical limitations that collectively hinder systematic and reproducible LALM evaluation, as summarized in Table 1. Despite varying in scope and design, no existing framework provides comprehensive coverage across the key dimensions required for rigorous evaluation: efficient and comprehensive inference backend integration, multi-turn dialogue support, concurrent multi-task execution, and configurable evaluation design. Together, these limitations force researchers to either compromise evaluation rigor or resort to fragmented, toolkit-specific workarounds that undermine reproducibility and cross-benchmark comparability. AU-Harness is the first unified and efficient toolkit to provide comprehensive support across all five dimensions, enabling systematic and reproducible LALM evaluation at scale. Our contributions are summarized as follows:

- We propose an **efficient evaluation engine** that leverages vLLM batching and dataset sharding to scale evaluations to multi-node infrastructures without sacrificing fidelity.
- A **unified, configurable framework** that standardizes evaluation protocols across benchmarks, enabling fair, reproducible comparisons and easy task integration.
- To the best of our knowledge, **AU-Harness** is the first unified evaluation framework to make **multi-turn** dialogue evaluation fully configurable, enabling systematic analyses of dialogue behaviors over extended interactions.

2 Related Work

Audio Benchmarks. Benchmarks play a critical role in the development of LALMs. SUPERB (Yang et al., 2021) established core task axes (Content, Speaker, Semantics, Paralinguistics) for audio model evaluation. DynamicSUPERB (Huang et al., 2024b) and DynamicSUPERB-2.0 (Huang et al., 2024a) expanded coverage to instruction-tuned and sequence generation tasks across speech, music, and environmental audio. Instruction-following and agentic conversational behaviors have been further probed by AIR-Bench (Yang et al., 2024) and VoiceBench (Chen et al., 2024). More recently, AudioBench (Wang et al., 2025a) unifies 8 task families over 26 datasets for LALMs.

Recent complementary efforts broaden the breadth and depth with audio reasoning capabilities: X-ARES (Zhang et al., 2025a) systematically assesses general audio encoders across domains, MECAT (Niu et al., 2025) targets fine-grained audio understanding with expert-guided captions and QA; MMAR (Ma et al., 2025), MMAU-PRO (Kumar et al., 2025), and MMSU (Wang et al., 2025b) focus on understanding and analyzing complex audio scenes, spatial relationships, and mixed-audio reasoning. Yet despite the rapid growth of audio benchmarks, the development of audio evaluation frameworks for fair and consistent comparisons between frontier models and benchmarks remains fairly understudied. This critical gap necessitates the development of a unified and efficient evaluation engine designed specifically for scalable audio evaluations.

Audio Evaluation Kits. In contrast with Audio Benchmark development, Audio Evaluation Kits have received less attention. This can be primarily attributed to the straightforward nature and minimal setup requirements of the early audio tasks, as presented in Dynamic-SUPERB-2.0 (Huang et al., 2024a) and AIR-Bench (Yang et al., 2024). How-

ever, the rapid growth of LALMs and the increasing complexity of newly curated audio benchmarks have underscored the critical need for comprehensive evaluation kits, as exemplified through the recent development of extensive evaluation kits (Ding et al., 2025; Wang et al., 2025a; Zhang et al., 2025b). For instance, AudioBench (Wang et al., 2025a) offers versatile evaluation support for up to 8 tasks across 26 datasets. VERSA (Shi et al., 2025) introduces a comprehensive framework to evaluate the quality of various speech, audio and music signals, with the focus on text-to-audio applications. AHELM (Lee et al., 2025) adds support for multi-turn audio conversation evaluation; however, its benchmark suite is largely fixed, and extending it to unsupported tasks requires writing task-specific code rather than specifying evaluations through a customizable configuration interface. Despite these advancements, most current evaluation kits operate on the simplified assumption that *a single model is evaluated against a single benchmark per run*. Addressing this limitation, we introduce an efficient, customizable evaluation kit to support the massive scale of the current LALMs and audio benchmarks as summarized in Table 1.

3 LALM Evaluation Challenges

3.1 Inference Efficiency

Most existing LALM evaluation kits have been designed based on the assumption that *a single model should be evaluated against a single benchmark per run*. However, this constrains researchers from conducting systematic, large-scale comparisons across LALMs and audio benchmarks efficiently, slowing the iterative process of model development and refinement. Current evaluation kits also under-utilize parallel processing capabilities available in the high-performance computing clusters, resulting in failures in incorporating benefits of available hardware infrastructures.

Two essential task-agnostic metrics for assessing the efficiency of LALM evaluation frameworks are *Real-time Factor (RTF)* and *Samples Processed per Second (SPS)*. RTF measures the processing time of an evaluation framework relative to the duration of the processed audio (Arriaga et al., 2024). Lower RTF is more desirable, indicating a more efficient audio evaluation framework. On the other hand, SPS directly quantifies the model’s processing speed by measuring the average number of audio samples processed per second. It serves as a complementary measure to RTF, providing a more

granular view of the model’s throughput and computational efficiency. The detailed formulation of these metrics are provided in Appendix A.4.

To quantify the efficiency of existing evaluation frameworks, we conduct a study on $N = 500$ audio samples (approximately 0.41 hours) of MELD-Emotion dataset (Poria et al., 2019) with experimental details in Appendix A.2. As reported in Table 1, most existing kits exhibit high RTF and low SPS, revealing fundamental throughput bottlenecks that ultimately hinder rapid LALM assessment over more diverse benchmark suites.

3.2 Customizable Evaluation Configurations

Multi-turn Dialogue Support Previous audio evaluation toolkits have largely been constrained to tasks centered on single-turn user interactions. However, as the field moves toward building interactive and context-aware voice assistants, the ability to evaluate multi-turn tasks becomes increasingly critical. Multi-turn evaluation enables a more realistic assessment of dialogue continuity, contextual reasoning, and the model’s capacity to adapt dynamically across extended conversations. Without such support, current evaluation approaches risk overlooking key aspects of usability and robustness that are essential for next-generation LALMs in realistic agentic voice systems.

Evaluation Customization. The lack of customizable filtering poses a significant barrier for researchers aiming to conduct in-depth analyses of current LALM limitations. Without the ability to refine evaluation datasets based on specific criteria, it is challenging to gain granular understanding of model performance across diverse audio conditions. For instance, while certain LALMs might perform reliably on 10-second audio chunks, they might be unable to handle short-form audio typically encountered in dialogue-state tracking systems.

4 AU-Harness

In response to the presented challenges, we propose a standardized, efficient, highly customizable evaluation framework, **AU-Harness**, detailed in Figure 1. AU-Harness is composed of 3 primary components: **Config**, **Central Request Controller (CRC)** and **Concurrent Engines**. The Config module defines a structured and hierarchical representation of customizable configurations, enabling flexible and transparent evaluation settings. The CRC is responsible for managing token requests and coordinating execution across the framework. Finally,




Figure 1: **Architecture overview of AU-Harness evaluation framework.** Our system comprises three core components: (1) *Config* module for hierarchical task configuration and standardized prompting, (2) *Central Request Controller (CRC)* managing token-based concurrency limits across all engines with adaptive retry mechanisms, and (3) *Concurrent Engines* executing parallel model evaluation with dataset sharding. The Central Request Controller maintains a global Token Pool accessible to all engines, enabling efficient resource utilization and scalable throughput. Multiple concurrent connections between the controller and inference models illustrate parallel request dispatch, with each engine supporting the evaluation of multiple models on targeted datasets.

the Concurrent Engines module carries out task-specific evaluations in parallel, where each engine can support multi-model evaluations tailored to particular tasks. In the following sections, we introduce our architecture design in detail to address the presented challenges in Section 3.

4.1 Inference Efficiency

As illustrated in Figure 1, AU-Harness maximizes inference throughput through three complementary mechanisms. First, we introduce a *Central Request Controller (CRC)* — a token-based scheduling architecture that maintains a global pool of concurrency slots shared across all models and evaluation engines. Each slot represents permission to issue one inference request; slots are acquired before dispatch and released upon completion, ensuring that throughput is governed solely by user-defined global request limits rather than model or engine-specific constraints. User-specified retry counts further provide a tunable balance between throughput and reliability, automatically re-attempting failed requests without manual intervention. Second, AU-Harness employs a *layered request synchronization* strategy that adaptively staggers wait times across concurrent models, increasing the probability that models processing the same dataset segment complete inference in a temporally aligned

manner; minimizing idle periods and intra-engine waiting time. Third, *proportional dataset sharding* partitions evaluation data into disjoint subsets distributed across model endpoints in proportion to each endpoint’s concurrency capacity, enabling improved throughput scaling across heterogeneous resources (as studied in Appendix A.5 and Figure 6). Together, these mechanisms deliver scalable, predictable evaluation throughput with minimal engineering overhead.

4.2 Customizable Evaluation Configurations

Backend-Agnostic Inference AU-Harness decouples predictive inference and metric computation from model hosting, requiring only a standardized model specification to integrate with the evaluation pipeline, regardless of whether the model is served via vLLM, a third-party API, or a custom FastAPI (Ramirez, 2018) endpoint. This design provides native support for vLLM-compatible models for high-throughput inference, while remaining fully compatible with any backend exposing a standard `/v1/chat/completions` interface. To lower the potential integration barriers, AU-Harness provides boilerplate FastAPI server implementations, allowing practitioners to wrap any optimized inference stack with minimal overhead and seamlessly integrate into the evaluation pipeline.

Evaluation customization. AU-Harness is also designed for granular control over evaluation steps. First, AU-Harness supports both open-source and proprietary models, which might contain their individualistic settings. Second, we allow for customizable metric assignment on a per-dataset and/or per-task basis. For instance, LLM-as-judge supports configurable concurrency to maximize the throughput for evaluation stage. For a more comprehensive understanding of model performance, the framework offers configurable aggregation metrics. This capability allows for the multi-dimensional analysis of task and metric results, providing a comprehensive outlook that extends beyond simple, individual scores or sub-tasks.

As shown in Figure 1, evaluation behavior is entirely specified through a YAML configuration, including the task-wise metrics, model configurations, optional score aggregation across related tasks, and prompt overrides. Importantly, no task-specific Python code or glue logic is required, even for complex benchmarks such as Emotion Recognition task suite that involves aggregation of multiple sub-datasets and LLM-based judging. This abstrac-




Figure 2: **Compositional task configuration in AU-Harness, illustrated on SpokenWOZ.** A shared base config defines all invariant components — processors, instruction prompt, generation parameters, and metrics — inherited by all task-specific variants. Each variant requires modifying only a single targeted field, such as input modality or transcription strategy, enabling new evaluation scenarios to be instantiated in minutes without any pipeline re-engineering. This compositional design is representative of AU-Harness’s broader configuration system, where comprehensive coverage of diverse tasks and settings is achieved through minimal, targeted modifications; facilitating rapid, reproducible experimentation at scale.

tion cleanly decouples tasks, metrics, models, and judges, enabling new evaluations to be launched by editing a single configuration file. User-friendly configuration is further illustrated in Appendix A.7.

4.3 Support for features and tasks

Multi-turn Dialogue support. Unlike existing toolkits, AU-Harness supports multi-turn evaluation in both audio and text modalities across LALMs through synchronous, turn-based evaluation chains that recursively append model outputs to the dialogue context at each turn. Beyond simple turn chaining, multi-turn audio evaluation introduces a non-trivial design space that existing toolkits leave unaddressed: from the second turn onward, the dialogue history may contain mixed-modality content, and the representation of prior turns must be explicitly configured. Specifically, each historical turn can be represented as (1) the original audio recording, (2) a ground-truth or reference text transcript, or (3) the model’s generated transcription from the previous pass. The optimal choice is not universal — it depends on the modality support and architectural constraints of the evaluated model. AU-Harness addresses this by providing dynamic configuration of history representation, enabling controlled and reproducible multi-turn evaluations across heterogeneous LALMs without requiring custom pipeline modifications for each model or dataset combination. This flexible design is further demonstrated via Section 5.2.

Continual engagement with emerging tasks. We aim to continue supporting up-and-coming benchmarks and models so AU-Harness can continue providing insights into LALM performance as they evolve. Figure 2 shows the ease of adding new datasets and evaluation configurations given the task category - the setup for a new task requires minimal updates with the given base configurations

to be supported in AU-Harness.

5 Results & Discussion

Without loss of generality, we adopt the task taxonomy proposed by Dynamic-SUPERB-2.0 (Huang et al., 2024a) for the empirical evaluations with our proposed AU-Harness due to its exhaustive coverage. Table 2 characterizes the breadth of audio evaluation suite supported by our AU-Harness, demonstrating the flexibility of our AU-Harness in supporting diverse audio tasks. Following (Wang et al., 2025a), we adopt GPT-4o-mini as judge for LLM-judge metrics³ due to its advanced capability. Further details of datasets, evaluated models and metrics are provided in Appendix A.1.

5.1 Inference Efficiency

Evaluation Settings. We benchmark AU-Harness against AudioBench (Wang et al., 2025a), VoiceBench (Chen et al., 2024), Kimi-Eval (Ding et al., 2025), and LMMS-Eval (Zhang et al., 2025b) on 500 audio samples across three diverse datasets and three LALMs, reporting averaged RTF and SPS. Additional runtime setups, namely *Sequential*, *Total* and *Parallel*⁴, to assure a comprehensive and fair comparison among all existing evaluation kits are also examined as detailed in Appendix A.2.

Evaluation Comparison. As shown in Figure 3, AU-Harness outperforms existing evaluation kits across different runtime scenarios in two key efficiency metrics. Against the most competitive baseline, LMMS-Eval with vLLM backend, AU-Harness achieves up to 77.2% relative reduction in RTFs under realistic end-to-end conditions (*Total* runtime). This improvement, as further demonstrated in Figure 4, reflects substantial throughput

³Results are obtained prior to the API discontinuation date

⁴Parallel refers to Parallel(Optimal) where no actual overhead is accounted for unless specified otherwise.

Models	Speech											Audio & Music	
Task Category	PR	ASR	Para.	S&L	SLU	HD	SE	Safety	Multi-turn		IF	AU	MU
Dataset	voxangeles	Librispeech	IEMOCAP	SR	BBA	StutterDetect	NoiseDetect	Advbench	MT-Bench	SpokenWoz	IFEval	AudioCaps	ChoMusic
Metrics	LB (↑)	WER (↓)	LB (↑)	LB (↑)	LBBA(↑)	LB (↑)	LB (↑)	SafetyJudge (↑)	MTJudge (↑)	JGA (↑)	IFScore (↑)	LB (↑)	LB (↑)
Small-sized Audio Language Models (<5B parameters)													
Voxtral-Mini-3B	0.10	2.10	54.90	45.80	43.50	12.90	14.50	78.50	65.88	24.92	40.02	14.96	45.40
Qwen2.5-Omni-3B	1.20	8.09	81.50	55.90	44.80	58.40	58.40	97.30	59.81	31.30	35.82	42.82	53.50
Medium Sized Large Audio Language Models (5B–20B parameters)													
Phi-4-Multi-modal	0.00	1.97	50.50	47.20	40.80	42.10	32.50	97.10	64.12	7.82	44.51	26.08	44.80
Qwen2.5-Omni-7B	21.60	1.74	85.80	62.30	50.90	68.20	59.00	98.30	64.56	37.30	50.83	38.40	59.30
Kimi-Audio	1.30	1.41	89.00	62.80	41.70	58.40	96.00	100.00	54.62	37.30	61.29	38.46	66.80
Large Sized Large Audio Language Models (>20B parameters)													
Voxtral-Small-24B	1.20	1.62	42.80	47.70	66.50	51.90	15.50	75.40	70.81	29.93	66.83	19.24	57.90
Qwen3-Omni-30B													
-A3B-Thinking	50.90	1.64	82.50	65.30	96.80	68.70	78.00	95.00	76.19	11.16	80.39	37.96	74.30
Proprietary Audio Language Models													
GPT-4o-mini-audio †	0.00	6.25	—*	40.30	63.70	3.80	53.00	88.10	65.00	28.39	70.47	15.08	50.20
Gemini2.5-Flash	35.80	2.17	92.70	60.20	90.30	64.60	87.50	98.50	74.50	52.09	84.63	36.16	72.90
Cascaded Systems													
Whisper-Large-v3 + GPT-oSS-20B	48.00	9.82	1.90	48.10	78.20	49.80	49.00	98.50	71.50	19.02	73.72	12.14	50.30
GPT-4o-transcribe + GPT-4.1-mini †	30.70	4.71	30.20	45.80	74.30	49.70	47.00	97.30	67.62	19.44	66.69	17.44	52.70

Table 2: **LALM performance on audio tasks.** We evaluate representative LALMs from different spectra: Open-source LALMs (small-sized, medium-sized, large-sized), Proprietary LALMs and Cascaded System LALMs across representative task categories. Metrics include LLM-as-judge evaluations using GPT-4o-mini and task-specific automatic metrics. **PR**: Phoneme Recognition, **Para.**: Paralinguistics, **ASR**: Automatic Speech Recognition, **S&L**: Speaker & Language, **SLU**: Spoken Language Understanding, **HD**: Hearing Disorder, **SE**: Speech Enhancement, **IF**: Instruction Following, **AU**: Audio Understanding, **MU**: Music Understanding, Refer to Appendix A.1 for benchmark, model and metric abbreviations together with their detailed explanations. **Bold**: highest; underline: second highest. *Performance affected by Azure OpenAI content filtering. † denotes results obtained prior to the API discontinuation date




Figure 3: **Efficiency comparison across evaluation frameworks and runtime scenarios.** (a) Real-time Factor (↓ better, log-scale) and (b) SPS (↑ better, log-scale) measured across three datasets (MELD-Emotion, LibriSpeech-test-clean, ClothoAQA) and three runtime conditions: Individual (dataset-specific), Sequential (worst-case serialized execution), Parallel (optimal concurrent execution) and Total (complete execution). AU-Harness consistently outperforms existing toolkits across all scenarios, with most significant gains in parallel and total execution, demonstrating effective utilization of concurrent processing.

gains when compared to contemporary evaluation frameworks across all evaluated datasets and models. The proposed orchestration layer (Figure 1) is designed to interface with any backend supporting dynamic scheduling, optimizing request handling and token delegation to minimize idle waiting time. As such, it is readily transferable to any inference backend sharing vLLM’s key characteristics: continuous batching, asynchronous request processing, and dynamic memory management.

Orchestration pipeline beyond vLLM integration. To validate the efficiency gains of CRC beyond naive vLLM integration, we conduct a comparative study against LMMS-Eval (Zhang et al., 2025b), a framework that similarly leverages vLLM for large-scale multimodal evaluation. The critical distinction lies in the scheduling strategy: while LMMS-Eval relies on vLLM’s default

processing pipeline with multi-threaded execution only, AU-Harness introduces dynamic concurrent request orchestration that actively minimizes idle waiting time across tasks. As shown in Table 3, although LMMS-Eval’s naive vLLM integration already improves throughput over HF-based counterparts (42.5% average SPS relative gain), it remains substantially less efficient than AU-Harness, which achieves an additional approx. 72.3% average SPS relative improvement. This gap widens further under realistic concurrent task processing conditions (*Reality*), where AU-Harness outperforms LMMS-Eval by 51.2% in RTF relative reduction and 109.2% in average SPS relative gain — demonstrating that effective vLLM utilization requires deliberate orchestration beyond default pipeline configurations.




Figure 4: **Total runtime efficiency analysis across evaluation frameworks.** Scatter plot comparing frameworks under optimal parallel execution conditions, plotting Real-time Factor (x-axis, ↓ better) against Samples Processed per Second (y-axis, ↑ better) on a log-log scale to span the wide dynamic range of both metrics (up to $\sim 63\times$ across frameworks). Our framework (top-left-most cluster) achieves superior performance in both dimensions at 8.79 s/s and 1.0% RTF, demonstrating the effectiveness of token-based request scheduling, dataset sharding, and vLLM integration for large-scale LALM evaluation.

EvalKit	RTF (↓)	SPS (↑)
LMMS-Eval-HF (Optimal)	3.11 ± 0.04	3.39 ± 0.12
LMMS-Eval-vLLM (Optimal)	1.87 ± 0.04	4.83 ± 0.10
LMMS-Eval-vLLM (Reality)	3.46 ± 0.22	2.62 ± 0.17
AU-Harness (Optimal)	<u>1.14 ± 0.16</u>	<u>8.32 ± 0.97</u>
AU-Harness (Reality)	<u>1.69 ± 0.08</u>	<u>5.48 ± 0.27</u>

Table 3: **Parallel runtime efficiency comparison between AU-Harness and LMMS-Eval.** We conduct controlled experiments following the setups presented in Section 5.1. As both LMMS-Eval and AU-Harness support multi-task parallel evaluation, besides the *Optimal* parallel runtime, we report *Reality* parallel runtime variant where additional realistic overheads are taken into account during evaluation. Clarifications are further provided in Appendix A.2. Bold and underline represent the best and second best performance respectively

5.2 Insights enabled by AU-Harness

A unified evaluation framework does more than evaluating models efficiently and comprehensively as it enables systematic analyses that would be difficult or impossible to conduct with fragmented toolkits. AU-Harness presents these through an interactive run report (Appendix A.8, Figures 8–14) covering category performance, error patterns, multi-turn dynamics, head-to-head comparison, and operational health. This enables further analyses including instruction modality effects, and multi-turn dialogue dynamics. Rather than an exhaustive survey, these case studies are intended to illustrate the analytical depth that a unified, extensible framework provides, and to highlight open challenges in the rigorous LALM evaluations.

Instruction Modality Gap. Table 4 reveals two compounding effects that expose fundamental limitations in current LALM evaluation practices. First, the modality of the instruction itself introduces measurable degradation: even under the strongest ASR condition (Whisper-v3, BERTScore 99.59), IFEval performance drops from 87.56 to 79.74

($\Delta = -7.82$), confirming that a residual instruction modality gap persists independent of transcription quality. Second, and more critically, the choice of evaluation paradigm dramatically shapes what benchmarks measure. On IFEval, the model’s own ASR achieves only 37.52 BERTScore and collapses task performance to 30.07 ($\Delta = -57.49$); yet even Qwen3-Captioner at 88.57 BERTScore only partially recovers performance to 58.76 ($\Delta = -28.80$), demonstrating that transcription quality alone does not predict task recovery. The pattern reverses on GSM8K: direct audio reasoning severely degrades performance ($\Delta = -30.93$), while Own ASR pre-processing recovers it to 72.02 ($\Delta = -5.31$) and Whisper-v3 fully closes the gap ($\Delta = +0.46$). This task-dependent asymmetry reveals that the cascaded intelligence bottleneck is not uniform: while instruction-following tasks are disproportionately sensitive to transcription artifacts, mathematical reasoning is primarily bottlenecked at the audio-grounded inference stage rather than the transcription stage. Uncovering these distinctions is feasible through a unified evaluation framework that supports modular pipelines and standardized evaluation protocols across diverse modality and inference conditions as offered in AU-Harness.

Multi-turn Dialogue Dynamics. Figure 5 presents JGA trajectories across six input-format configurations on SpokenWOZ, each enabled by AU-Harness’s flexible multi-turn evaluation design. Across all standard configurations, JGA declines sharply with each dialogue turn and approaches zero by turn 14. This trend closely tracks the monotonically increasing number of accumulated slots to predict, indicating that cumulative slot

Benchmark	Text base.	Prompt-only				← ASR Preprocessing (increasing quality) →					
		Direct audio	Δ	2-step prompt	Δ	Own ASR	Δ	Qwen3 Cap. [†]	Δ	Whisper -v3	Δ
INSTRUCTION FOLLOWING											
IFEval	87.56	80.39	-7.17	71.17	-16.39	30.07	-57.49	58.76	-28.80	79.74	-7.82
MT-Bench	89.88	82.56	-7.32	81.94	-7.94	61.88	-28.00	89.44	-0.44	88.38	-1.50
MATHEMATICAL & COMPLEX REASONING											
GSM8K	77.33	46.40	-30.93	63.31	-14.02	72.02	-5.31	76.72	-0.61	77.79	+0.46
BBH	97.50	92.30	-5.20	94.80	-2.70	89.49	-8.01	94.90	-2.60	97.40	-0.10

Benchmark	BERTScore ↑		
	Own ASR	Qwen3 Cap. [†]	Whisper -v3
INSTR. FOLLOWING			
IFEval	37.52	88.57	99.59
MT-Bench	75.59	94.54	95.98
REASONING			
GSM8K	92.79	97.94	98.37
BBH	90.61	94.46	95.35

[†] Qwen3-Captioner: Qwen3-family ASR-specialized model, only applicable for Qwen3-Omni. Color coding (BERTScore): green ≥ 95 , amber 70–94, red < 70 . Red background: $|\Delta| > 10$ pts (severe degradation). Green background: $|\Delta| \leq 10$ pts (minor degradation or improvement).

Table 4: **Instruction Modality Gap — Qwen3-Omni-30B-A3B-Thinking** *Left*: Task performance across benchmarks under different evaluation paradigms. Δ = Score – Text baseline; negative values indicate degradation. Red: $|\Delta| > 10$ pts. Green: $|\Delta| \leq 10$ pts. ASR preprocessing columns ordered by increasing transcription quality. *Right*: ASR transcription quality (BERTScore $\uparrow$) per benchmark $\times$ ASR condition. Detailed experimental settings are provided in Appendix A.6.




Figure 5: **SpokenWOZ multi-turn DST — JGA vs. turn index across input-format configurations for Qwen3-Omni-30B-A3B-Thinking.** Each panel groups evaluation configurations enabled by AU-Harness by scoring strategy: *Standard Configurations* (left) vary the current-turn modality across text-only (baseline), last user turn in audio, and all user turns in audio, evaluated against the cumulative dialogue state; *Standard Configurations with GT Prior State* (right) apply the same modality variants but inject the oracle prior state before each turn, scoring only the per-turn delta. Grey bars denote the mean accumulated slot count per turn (left panel) and the mean number of new slots introduced per turn (right panel), both on the right axis. Further details in Section A.9.

complexity is a major driver of performance collapse. The oracle configurations isolate this effect by providing the ground-truth prior state and scoring only per-turn slot deltas, thereby removing accumulated slot errors as a confounder. Under this setting, JGA stabilizes substantially, suggesting that unfilled slot accumulation, not turn-level understanding, is the dominant bottleneck. The oracle curves therefore provide a more faithful estimate of per-turn comprehension capacity. Text-only inputs achieve the highest JGA in both standard and oracle settings, consistent with strong transfer from extensive text-based pretraining in LALMs. Presenting the current turn in audio lowers JGA relative to the text-only baseline, revealing a persistent audio-text perception gap even when only the current turn is spoken. This gap becomes more pronounced under the oracle setting, suggesting that it is partially masked by error accumulation in standard cumulative-state evaluation. The all-audio configuration, in which all user turns are spoken, exhibits an ambiguous trend under standard evaluation but shows substantially lower JGA under the oracle

setting, exposing a more fundamental audio-only comprehension gap. Crucially, this analysis spans modality, history representation, and state-tracking strategy within a single evaluation framework. AU-Harness’s composable, per-model multi-turn configuration system enables these comparisons directly, whereas existing toolkits would require bespoke pipeline engineering.

6 Conclusion

We introduced a modular and extensible evaluation framework for large audio-language models that emphasizes broad task coverage, ease of use, and adaptability. Its modular design enables researchers and practitioners to extend the codebase, customize benchmarks, and integrate new models or tasks without major restructuring. The efficiency gains of our AU-Harness are realized through the aggregation of dataset sharding and effective token request orchestration. More importantly, the broader value of our framework lies in enabling flexible, large-scale evaluations that were previously difficult to conduct in a reproducible and accessible manner. By lowering the barrier to bench-

marking and fostering customization, we aim to support both systematic research and practical deployment, contributing a more standardized and transparent evaluation ecosystem for LALMs.

Limitations

Backend dependency and reproducibility. Our efficiency gains are evaluated with vLLM integration; hence, models without mature backends might revert to conventional execution with reduced throughput. Support for closed-source endpoints depends on chat-completions APIs, limiting batching control and introducing provider rate limits. Even with deterministic configs, runs may vary due to endpoint queuing and transient failures, requiring documentation of capacity and request budgets for cross-institutional comparability.

Standardization vs. task fidelity. Standardized prompting improves reproducibility but cannot eliminate prompt sensitivity. For open-ended tasks, canonical prompts may bias results toward specific behaviors. The community needs multiple documented prompt families and complementary temporal measures to triangulate performance fairly.

Evaluation Framework vs Benchmark. Our work presents a unified and efficient evaluation framework targeting specifically for LALM evaluation. Despite the wide coverage presented in Table 2, our ultimate objective is to create an extensible framework where new benchmarks, tasks and metrics can be seamlessly integrated.

Coverage and generalization gaps. Our coverage remains skewed toward English and common domains. Environmental audio, music understanding, and low-resource languages are underrepresented. Moreover, the relationship between standardized benchmark performance and real-world audio-language capabilities where contexts are noisier, more diverse, and less structured requires further empirical validation.

Limited analytical study scope Both insight analyses presented in Section 5.2 are intentionally scoped as representative case studies: the instruction modality and cascaded intelligence study examines only Qwen3-Omni-30B-A3B-Thinking, and the multi-turn DST study evaluates one model on one benchmark (SpokenWOZ) across diverse input configurations. Whether the observed modality gaps and turn-level degradation generalize across broader LALM families remains an open question for future investigations. These analyses are not intended to be exhaustive; instead, they demon-

strate the range and depth of systematic investigations that AU-Harness enables. Ultimately, AU-Harness aims to provide a standardized, efficient, highly customizable evaluation framework for more comprehensive studies by the broader research community.

These limitations highlight challenges in audio-language evaluation. Achieving reproducible, comprehensive, and valid assessment requires community coordination around evaluation standards, temporal diagnostics, and multilingual breadth. Our framework is designed to enable practical, systematic progress in these areas across the broader ecosystem.

Acknowledgments

We extend our gratitude to the CLAE team at ServiceNow for their invaluable feedback on the architecture design of our evaluation framework.

Ethics Statement

Our work focuses on responsible development of audio language model evaluation infrastructure. We have taken care to ensure that all audio datasets used in our benchmarks respect copyright and privacy guidelines, with particular attention to speaker consent in diarization tasks. We also strictly adhere to the official documentation and intended use protocols of all evaluated frameworks and benchmarks accordingly. While our framework enables large-scale evaluation of LALMs, we cannot guarantee that models evaluated through AU-Harness will not generate harmful or biased audio-related outputs. Researchers and practitioners are strongly encouraged to implement appropriate content filtering and bias detection when deploying LALMs in production environments. Our speech synthesis components for creating reasoning benchmarks use only publicly available, ethically sourced voice models. Additionally, we acknowledge that our current task coverage is skewed toward English and common domains, which may inadvertently reinforce existing representational biases in audio AI systems. We encourage the community to extend our framework to include more diverse languages and cultural contexts.

Regarding language model usage in manuscript preparation, we utilize them solely to refine the language used in paper to improve clarity and correctness, without generating any substantial content or claims.

Reproducibility Statement

We are committed to full reproducibility of our evaluation framework and experimental results. All AU-Harness code, configuration files, evaluation scripts, and documentation will be publicly released under an open-source license upon acceptance. We provide comprehensive implementation details including all hyperparameters, model endpoints, dataset preprocessing steps, and evaluation metrics in our appendices. For efficiency comparisons, we document exact hardware specifications, vLLM versions, concurrent request limits, and retry policies used across all experiments. Our newly introduced reasoning benchmarks include complete details on text-to-speech synthesis parameters and prompt templates. To ensure consistent reproduction, we provide Docker containers with fixed dependency versions and detailed setup instructions for multi-node evaluation. All random seeds, sampling parameters, and LLM-as-judge configurations are specified to enable identical result replication across different research groups.

References

- Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, and 1 others. 2025. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*.
- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, and 1 others. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Carlos Arriaga, Alejandro Pozo, Javier Conde, and Alvaro Alonso. 2024. Evaluation of real-time transcriptions using end-to-end asr models. *arXiv preprint arXiv:2409.05674*.
- Yiming Chen, Xianghu Yue, Chen Zhang, Xiaoxue Gao, Robby T Tan, and Haizhou Li. 2024. Voicebench: Benchmarking llm-based voice assistants. *arXiv preprint arXiv:2410.17196*.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Santiago Cuervo, Skyler Seto, Maureen de Seyssel, Richard Bai, Zijin Gu, Tatiana Likhomanenko, Navdeep Jaitly, and Zakaria Aldeneh. 2026. Closing the gap between text and speech understanding in llms. In *International Conference on Learning Representations*, volume 2026, pages 146313–146335.
- Wenqian Cui, Dianzhi Yu, Xiaoqi Jiao, Ziqiao Meng, Guangyan Zhang, Qichao Wang, Yiwen Guo, and Irwin King. 2024. Recent advances in speech language models: A survey. *arXiv preprint arXiv:2410.03751*.
- Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song, Xu Tan, Heyi Tang, and 1 others. 2025. Kimi-audio technical report. *arXiv preprint arXiv:2504.18425*.
- Chien-yu Huang, Wei-Chih Chen, Shu-wen Yang, Andy T Liu, Chen-An Li, Yu-Xiang Lin, Wei-Cheng Tseng, Anuj Diwan, Yi-Jen Shih, Jiatong Shi, and 1 others. 2024a. Dynamic-superb phase-2: A collaboratively expanding benchmark for measuring the capabilities of spoken language models with 180 tasks. *arXiv preprint arXiv:2411.05361*.
- Chien-yu Huang, Ke-Han Lu, Shih-Heng Wang, Chi-Yuan Hsiao, Chun-Yi Kuan, Haibin Wu, Siddhant Arora, Kai-Wei Chang, Jiatong Shi, Yifan Peng, and 1 others. 2024b. Dynamic-superb: Towards a dynamic, collaborative, and comprehensive instruction-tuning benchmark for speech. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12136–12140. IEEE.
- Ivan Ismanto, Indra Gamayanto, and Gabriello Klavin Sanyoto. 2026. Improvement of user experience evaluation for smes digital application using tri, tam, sus integration. *Journal of Applied Informatics and Computing*, 10(1):877–888.
- Ishika Joshi, Simra Shahid, Shreeya Manasvi Venneti, Manushree Vasu, Yantao Zheng, Yunyao Li, Balaji Krishnamurthy, and Gromit Yeuk-Yin Chan. 2025. Coprompter: User-centric evaluation of llm instruction alignment for improved prompt engineering. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*, pages 341–365.
- Sonal Kumar, Šimon Sedláček, Vaibhavi Lokegaonkar, Fernando López, Wenyi Yu, Nishit Anand, Hyeon-gon 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. *arXiv preprint arXiv:2508.13992*.
- Sonal Kumar, Šimon Sedláček, Vaibhavi Lokegaonkar, Fernando López, Wenyi Yu, Nishit Anand, Hyeon-gon Ryu, Lichang Chen, Maxim Plička, Miroslav Hlaváček, and 1 others. 2026. Mmau-pro: A challenging and comprehensive benchmark for holistic evaluation of audio general intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 22688–22697.
- Tony Lee, Haoqin Tu, Chi Heem Wong, Zijun Wang, Siwei Yang, Yifan Mai, Yuyin Zhou, Cihang Xie,

- and Percy Liang. 2025. Ahelm: A holistic evaluation of audio-language models. *arXiv preprint arXiv:2508.21376*.
- Alexander H Liu, Andy Ehrenberg, Andy Lo, Clément Denoix, Corentin Barreau, Guillaume Lample, Jean-Malo Delignon, Khyathi Raghavi Chandu, Patrick von Platen, Pavankumar Reddy Muddireddy, and 1 others. 2025. Voxtral. *arXiv preprint arXiv:2507.13264*.
- 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. *arXiv preprint arXiv:2505.13032*.
- Dena Mujtaba, Nihar Mahapatra, Megan Arney, J Yaruss, Hope Gerlach-Houck, Caryn Herring, and Jia Bin. 2024. Lost in transcription: Identifying and quantifying the accuracy biases of automatic speech recognition systems against disfluent speech. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4795–4809.
- Hoang Nguyen, Chenwei Zhang, Ye Liu, Natalie Parde, Eugene Rohrbaugh, and Philip S Yu. 2024. Cori: Cjkv benchmark with romanization integration-a step towards cross-lingual transfer beyond textual scripts. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4008–4020.
- Hoang Nguyen, Chenwei Zhang, Tao Zhang, Eugene Rohrbaugh, and Philip S Yu. 2023. Enhancing cross-lingual transfer via phonemic transcription integration. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9163–9175.
- Hoang H Nguyen, Khyati Mahajan, Vikas Yadav, Julian Salazar, Philip S Yu, Masoud Hashemi, and Rishabh Maheshwary. 2025. Prompting with phonemes: Enhancing llms’ multilinguality for non-latin script languages. In *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)*, pages 11975–11994.
- Yadong Niu, Tianzi Wang, Heinrich Dinkel, Xingwei Sun, Jiahao Zhou, Gang Li, Jizhong Liu, Xunying Liu, Junbo Zhang, and Jian Luan. 2025. Mecat: A multi-experts constructed benchmark for fine-grained audio understanding tasks. *Preprint*, arXiv:2507.23511.
- OpenAI. 2024. GPT-4o mini Audio (preview). <https://platform.openai.com/docs/models/gpt-4o-mini-audio-preview>. Model: gpt-4o-mini-audio-preview-2024-12-17.
- OpenAI. 2025a. GPT-4o Transcribe. <https://platform.openai.com/docs/models/gpt-4o-transcribe>. Model: gpt-4o-transcribe; added to the Audio API alongside gpt-4o-mini-transcribe.
- OpenAI. 2025b. Introducing GPT-4.1 in the API. <https://openai.com/index/gpt-4-1/>. Model: gpt-4.1-mini-2025-04-14.
- Shishir G. Patil, Huanzhi Mao, Charlie Cheng-Jie Ji, Fanjia Yan, Vishnu Suresh, Ion Stoica, and Joseph E. Gonzalez. 2024. The berkeley function calling leaderboard (bfcl): From tool use to agentic evaluation of large language models. In *Advances in Neural Information Processing Systems*.
- Bornali Phukon, Xiuwen Zheng, and Mark Hasegawa-Johnson. 2025. Aligning ASR Evaluation with Human and LLM Judgments: Intelligibility Metrics Using Phonetic, Semantic, and NLI Approaches. In *Interspeech 2025*, pages 5708–5712.
- 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 *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 527–536.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pages 28492–28518. PMLR.
- Sebastián Ramirez. 2018. Fastapi. <https://fastapi.tiangolo.com/>.
- Jiatong Shi, Hye-jin Shim, Jinchuan Tian, Siddhant Arora, Haibin Wu, Darius Petermann, Jia Qi Yip, You Zhang, Yuxun Tang, Wangyou Zhang, and 1 others. 2025. Versa: A versatile evaluation toolkit for speech, audio, and music. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, pages 191–209.
- 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. *arXiv preprint arXiv:2310.13289*.
- Bin Wang, Xunlong Zou, Geyu Lin, Shuo Sun, Zhuohan Liu, Wenyu Zhang, Zhengyuan Liu, AiTi Aw, and Nancy F. Chen. 2025a. AudioBench: A universal benchmark for audio large language models. In *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)*, pages 4297–4316, Albuquerque, New Mexico. Association for Computational Linguistics.

- Dingdong Wang, Jincenzi Wu, Junan Li, Dongchao Yang, Xueyuan Chen, Tianhua Zhang, and Helen Meng. 2025b. Mmsu: A massive multi-task spoken language understanding and reasoning benchmark. *arXiv preprint arXiv:2506.04779*.
- Bajian Xiang, Shuaijiang Zhao, Tingwei Guo, and Wei Zou. 2025. [Understanding the modality gap: An empirical study on the speech-text alignment mechanism of large speech language models](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5187–5202, Suzhou, China. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv:2503.20215*.
- Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfa Zhu, and 1 others. 2025b. Qwen3-omni technical report. *arXiv preprint arXiv:2509.17765*.
- Qian Yang, Jin Xu, Wenrui Liu, Yunfei Chu, Ziyue Jiang, Xiaohuan Zhou, Yichong Leng, Yuanjun Lv, Zhou Zhao, Chang Zhou, and 1 others. 2024. Airbench: Benchmarking large audio-language models via generative comprehension. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1979–1998.
- Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhota, Yist Y Lin, Andy T Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, and 1 others. 2021. Superb: Speech processing universal performance benchmark. *Interspeech 2021*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- Junbo Zhang, Heinrich Dinkel, Yadong Niu, Chenyu Liu, Si Cheng, Anbei Zhao, and Jian Luan. 2025a. [X-ares: A comprehensive framework for assessing audio encoder performance](#). *Preprint*, arXiv:2505.16369.
- Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and 1 others. 2025b. Lmms-eval: Reality check on the evaluation of large multimodal models. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 881–916.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

A Appendix

A.1 Comprehensive Audio Evaluation

A.1.1 Benchmark Details

We present a comprehensive benchmark suite comprising 61 diverse datasets spanning six fundamental task categories in audio and speech understanding as observed in Table 5. Our benchmark suite encompasses *Audio Understanding* (6 datasets), evaluating models’ capabilities in audio scene analysis and music comprehension; *Paralinguistics* (12 datasets), assessing speech characteristics including emotion, gender, accent recognition, and speaker-related tasks; *Safety and Security* (2 datasets), examining robustness against adversarial inputs and spoofing; *Spoken Language Reasoning* (5 datasets), testing complex reasoning abilities from mathematical problem-solving to code generation from speech; *Spoken Language Understanding* (21 datasets), the largest category covering speech question-answering, intent classification, and translation tasks; and *Speech Recognition* (15 datasets), establishing baselines for automatic speech recognition across multiple languages and acoustic conditions. The licenses of individual datasets are also further noted in Table 5. Evaluating individual tasks with AU-Harness follows the licensing terms as denoted in Appendix A.10.

Spoken Language Reasoning (SLR) Existing SLR benchmarks are sparse and lack rigorous quality validation. For instance, VoiceBench (Chen et al., 2024) covers a multi-turn audio reasoning task but covers only a subset of the equivalent text-based multi-turn reasoning task (MT-Bench (Zheng et al., 2023)). To address this gap, we construct targeted SLR tasks by systematically converting challenging text-based reasoning benchmarks into the audio form via high-quality Text-to-Speech (TTS) systems. Our work centers on 3 major reasoning tasks, including:

- **Speech Function Calling:** Assesses the LALMs' comprehension of spoken instructions and their ability to map spoken natural language queries into structured, executable function calls with appropriate arguments. We achieve the goal by expanding BFCL-v3 (Patil et al., 2024) by systematically converting textual instructions into spoken counterparts.
- **Speech-to-Coding:** Evaluates LALMs' capability to translate spoken instructions into a formal programming language. Adapted from the renowned Spider text-to-SQL benchmark (Yu et al., 2018), we construct Speech-Spider benchmark where LALMs are expected to convert spoken instructions into valid SQL commands.
- **Speech Instruction Following (Speech-IF):** Evaluates proficiency in interpreting and executing intricate, potentially multi-step audio instructions. We develop Speech-MTBench benchmarks, derived from the well-known text-based MTBench (Zheng et al., 2023) to assess this capability.

A.1.2 Metric Details

- **Word Error Rate (WER)** – Measures automatic speech recognition (ASR) errors via insertions and deletions in transcribed text. Lower is better.
- **LLM-Judge (MJ)** – LLM-based evaluation of response quality. Higher is better. Reported metrics:
 - **Binary (LB)** – Binary LLM-based pass/fail correctness judgment.
 - **Detailed (LD)** – Detailed multi-level llm judgment across multiple dimensions.
 - **BigBench Audio (LBBA)** – LLM-based evaluations for BigBench-like audio tasks.
 - **RedTeaming (SafetyJudge)** – LLM-based evaluations for red-teaming and safety.
 - **MT-Bench (MTJudge)** – LLM-based evaluation for multi-turn systems.
- **BLEU** – N-gram overlap score for comparing generated and reference text. Higher is better.

¹https://huggingface.co/datasets/ServiceNow-AI/BFCL_v3_audio

²https://huggingface.co/datasets/ServiceNow-AI/mt-bench_audio

³https://huggingface.co/datasets/ServiceNow-AI/spider_audio

- **Instruction Following Score (IFScore)** (Zhou et al., 2023) – Measuring instruction following capability in natural language tasks via averaging accuracy across (1) strict-prompt, (2) strict-instruction, (3) loose-prompt and (4) loose-instruction scenarios.
- **Joint Goal Accuracy (JGA)** - A strict holistic measure for dialogue state tracking systems requiring a perfect alignment between predicted states and reference states for every slot-value pair at every single turn in multi-turn conversations.

A.1.3 Evaluated Models

We evaluate a diverse set of LALMs spanning open-source and proprietary systems, organized by model size and architecture. A brief description of each evaluated model is provided below.

Small-sized LALMs (<5B parameters)

Voxtral-Mini-3B (Liu et al., 2025)

A compact, open-source audio language model developed by Mistral AI, designed for efficient speech understanding and instruction following at a small parameter budget.

Qwen2.5-Omni-3B (Xu et al., 2025a)

The smallest variant of the Qwen2.5-Omni model family, an omni-modal architecture from Alibaba that jointly processes audio, vision, and text within a unified framework.

Medium-sized LALMs (5B–20B parameters)

Phi-4-Multimodal (Abouelenin et al., 2025)

A multimodal model developed by Microsoft, extending the Phi-4 language model with speech and audio perception capabilities through a lightweight LoRA architecture design.

Qwen2.5-Omni-7B (Xu et al., 2025a)

The standard-scale variant of the Qwen2.5-Omni family, offering a strong balance between audio understanding capability and computational efficiency.

Kimi-Audio (Ding et al., 2025)

An open-source audio language model developed by Moonshot AI, trained on large-scale audio-text paired data with particular strength in ASR and speech enhancement tasks.

Large-sized LALMs (>20B parameters)

Voxtral-Small-24B (Liu et al., 2025)

The larger variant of the Voxtral series from Mistral AI, providing substantially improved instruction-

Task Category	Task Type	Dataset Name	Task	Description	License
Speech Recognition	ASR	AISHELL-1	1	High-quality Mandarin speech recognition dataset	Apache 2.0
	ASR	AMI Meeting Corpus	2	Multispeaker meeting recordings for ASR and diarization	CC BY 4.0
	ASR	CallHome	5	Conversational speech corpus across multiple languages	LDC User Agreement for Non-Members
	ASR	Common Voice	100	Crowdsourced multilingual speech dataset from Mozilla	CC0 1.0 Universal
	ASR	FLEURS EN-US	102	Multilingual speech dataset for ASR and translation	CC BY 4.0
	ASR	GigaSpeech	1	Large-scale audio and transcription corpus for end-to-end ASR	Apache 2.0
	ASR	GigaSpeech2	2	Large-scale audio and transcription corpus for end-to-end ASR (v2)	Apache 2.0
	ASR	LibriSpeech	2	Audiobook-derived speech corpus with clean and noisy subsets	CC BY 4.0
	ASR	Multilingual LibriSpeech (MLS)	7	Extension of LibriSpeech with multiple European languages	CC BY 4.0
	ASR	MNSC	6	Large-scale multitask speech corpus	MNSC: Publicly released
	ASR	People’s Speech	1	Large-scale open-source English speech recognition dataset	CC-BY-SA
	ASR	SPGISpeech	1	Transcriptions of financial meeting recordings	Kensho User Agreement
	ASR	TEDLIUM3	1	Transcribed TED talks for ASR and speaker adaptation	CC BY-NC-ND 3.0
	ASR	VoxPopuli	17	Multilingual speech corpus from European Parliament recordings	CC0
	Code-switching ASR	SEAME	2	Mandarin-English code-switching speech dataset	LDC2015S04
	Long-form ASR	TEDLIUM3	1	Transcribed TED talks for ASR and speaker adaptation (long-form version)	CC BY-NC-ND 3.0
	Long-form ASR	Earnings21	1	Long-form earnings call dataset for speech recognition	CC-BY-SA-4.0
	Long-form ASR	Earnings22	1	Long-form earnings call dataset for speech recognition	CC-BY-SA-4.0
Paralinguistics	Accent Recognition	MNSC AR Dialogue	1	Dataset for accent recognition in dialogue speech	MNSC: Publicly released
	Accent Recognition	MNSC AR Sentence	1	Dataset for accent recognition in sentence-level speech	MNSC: Publicly released
	Accent Recognition	VoxCeleb Accent	1	Speech dataset with diverse speakers for accent recognition	CC BY 4.0
	Emotion Recognition	IEMOCAP Emotion	1	Multimodal dataset for emotion recognition in speech	GPL-3.0
	Emotion Recognition	MELD Emotion	1	Multi-party conversation dataset for emotion recognition	GPL-3.0
	Emotion Recognition	MELD Sentiment	1	Multi-party conversation dataset for sentiment analysis	GPL-3.0
	Gender Recognition	IEMOCAP Gender	1	Multimodal dataset for gender recognition in speech	GPL-3.0
	Gender Recognition	MNSC GR Dialogue	1	Dataset for gender recognition in dialogue speech	MNSC: Publicly released
	Gender Recognition	MNSC GR Sentence	1	Dataset for gender recognition in sentence-level speech	MNSC: Publicly released
	Gender Recognition	VoxCeleb Gender	1	Speech dataset with diverse speakers for gender recognition	CC BY 4.0
Audio Understanding	Speaker Recognition	MMAU-mini	1	Multi-modal audio dataset for speaker recognition	Apache 2.0
	Music Understanding	MuChocoMusic	1	Benchmark for music understanding for LALMs	CC-BY-SA-4.0
	Scene Understanding	AudioCaps	1	Large-scale dataset for open-domain audio captioning	MIT
	Scene Understanding	AudioCaps QA	1	Dataset for question answering over natural audio scenes	MIT
	Scene Understanding	Clotho AQA	1	Dataset for answering natural-language questions about audio signals	MIT
	Scene Understanding	WavCaps	1	Large-scale weakly labeled dataset for audio captioning	CC-BY-NC 4.0
	Scene Understanding	WavCaps QA	1	Large-scale dataset for audio question answering	CC-BY-NC 4.0
Spoken Language Understanding	Intent Classification	SLURP	1	Multi-domain spoken dialogue understanding benchmark	CC BY-NC 4.0
	Speech QA	Alpaca Audio	1	Speech dataset for question answering with audio instructions	Apache-2.0
	Speech QA	CN College Listen MCQ	1	Multispeaker dataset for listening-based multiple-choice questions	MERaLiON Public License
	Speech QA	Dream TTS MCQ	1	Dialogue-based multiple-choice comprehension dataset with audio	MIT
	Speech QA	MNSC SQA	4	Benchmark for reasoning and understanding in spoken language	NSC License
	Speech QA	OpenHermes	1	Speech dataset for question answering with audio instructions	CC-BY-NC
	Speech QA	Public-SG	1	Speech question answering benchmark	NSC License
	Speech QA	SLUE SQA	1	Spoken Language Understanding Evaluation benchmark	CC-BY-4.0
	Speech QA	Spoken Squad	1	Speech dataset for extraction-based question answering	CC-BY-SA-4.0
	SQQA	Big Bench Audio	2	Benchmark for reasoning with audio and text input	MIT
	SQQA	MMSU	12	Multi-choice question answering dataset	Apache-2.0
	SQQA	OpenBookQA	1	Multi-choice question answering dataset	Apache-2.0
	SQQA	SD-QA	22	Multi-choice question answering dataset	Apache-2.0
	Translation	CoVoST2 (zh→en)	36	Large-scale multilingual dataset for speech translation	CC-BY-NC-4.0
Spoken Language Reasoning	Speech Instruction Following	IFEVAL	2	Speech and text-equivalent dataset for complex instruction following	Apache-2.0
	Speech Instruction Following	VoiceBench-MTBench (46/80 samples)	1	Speech dataset for multi-turn instruction following	Apache-2.0
	Grade School Math	GSM8k	2	Speech and text-equivalent dataset of grade school math word problems	MIT (text dataset)
	Speech Function Calling	BFCL ¹	20	Speech and text-equivalent dataset for complex function calling tasks	Apache-2.0
	Speech Instruction Following	MTBench ²	2	Speech and text-equivalent dataset for multi-turn instruction following	Apache-2.0
	Speech-to-Coding	SPEECH_TO_SQL ³	2	Speech and text-equivalent dataset for generating executable SQL code	Apache-2.0
Safety & Security	Safety	Advbench	1	Speech dataset for testing resistance to adversarial or harmful prompts	Apache 2.0
	Spooing	ASVpoof2017	1	Speech dataset for spoofing attack detection in real-world conditions	CC BY-NC 4.0
Total Tasks			387		

Table 5: **Comprehensive Audio and Speech Datasets Overview.** Listing of 61 datasets across 6 task categories: Speech Recognition, Paralinguistics, Audio Understanding, Spoken Language Understanding, Spoken Language Reasoning, and Safety & Security.

following and spoken language understanding capabilities over its 3B counterpart.

Qwen3-Omni-30B-A3B-Thinking (Xu et al., 2025b)

A large-scale mixture-of-experts omni-modal model from Alibaba, operating with 3B active parameters out of 30B total. It incorporates an explicit thinking mode that emits internal reasoning traces prior to generating a final response, targeting complex reasoning and instruction-following tasks.

Proprietary LALMs

GPT-4o-mini-audio (OpenAI, 2024)

A cost-efficient variant of OpenAI’s GPT-4o model with native audio input and output capabilities, supporting real-time spoken interaction and multi-modal understanding.

Gemini-2.5-Flash (Comanici et al., 2025)

State-of-the-art proprietary multimodal model with strong audio, vision, and text understanding, optimized for low-latency inference while maintaining

high performance across diverse tasks.

Cascaded Systems

Whisper-Large-v3 + GPT-oSS-20B (Radford et al., 2023; Agarwal et al., 2025)

A cascaded pipeline combining OpenAI’s state-of-the-art ASR model Whisper-Large-v3 for transcription with GPT-oSS-20B for downstream language understanding and reasoning, representing a strong ASR-then-LLM baseline.

GPT-4o-transcribe + GPT-4.1-mini (OpenAI, 2025a,b)

A cascaded pipeline pairing OpenAI’s GPT-4o-based transcription model with GPT-4.1-mini for task completion, providing a fully proprietary transcription-then-reasoning baseline.

A.2 Inference Efficiency Evaluation Settings

To provide a comprehensive and fair comparison with other evaluation kits, regardless of their underlying implementation, we introduce two addi-

tional runtime scenarios beyond individual dataset runtimes, namely *Sequential* and *Parallel*. First, *Sequential* runtime represents the most inefficient runtime by assuming each benchmark is executed in a sequential manner, where no data or model parallelization algorithms are introduced. On the other hand, *Parallel* presents the theoretical upper-bound for optimal runtime. The final *Parallel* runtime is calculated by taking the longest runtime among all evaluated datasets. This scenario presumes an ideal, zero-overhead parallelization environment where communication protocols among parallel processes and other overheads do not impact the runtime. This is considered a best-case runtime for our framework and existing evaluation kits across all presented datasets and models. On the other hand, we also define *Parallel (Reality)* as an end-to-end runtime metric when all system overheads are taken into consideration. This metric is only reported for evaluation frameworks that natively support concurrent task evaluation. Evaluation frameworks lacking these capabilities do not have *Parallel (Reality)* runtime.

In our experimental settings, for fair comparison, we allocate 3xH100 GPUs to all of the evaluation frameworks and maximize the throughput designed by the frameworks either through multi-processing or supported concurrent parallel multi-task evaluations. The *Total* runtime (reported in Figure 3, Figure 4 and Section 5) reflects evaluation of all models and tasks under each framework’s maximum supported conditions. To further ensure fair comparison, caching mechanisms are disabled across all frameworks, preventing any artificial throughput gains from repeated inference runs on previously seen inputs.

For LLM judge, we utilize GPT-4.1-mini, queried via Azure OpenAI API support. The choice of GPT models helps eliminate inference overhead from local model hosting. For all frameworks, the number of concurrency requests (or equivalent batch size, depending on the framework’s inference design) is fixed at 20 and the judge concurrency is similarly set to 20. The complete evaluated settings are provided in Table 6. Our reported results represent the mean and standard deviations computed over 3 independent runs unless stated otherwise. This consistency is applied for Section 3.1 and Section 5.1.

A.3 Contemporary Evaluation Kits

There are a few evaluation kits that we have built upon and been inspired by, both in evaluation framework design and task coverage.

- **AudioBench** (Wang et al., 2025a): A comprehensive open-source audio evaluation framework encompassing eight core tasks and more than twenty-six curated datasets, with coverage continuing to expand. AudioBench supports both open and closed-source models and provides standardized evaluation pipelines using conventional metrics such as Word Error Rate (WER) and METEOR, alongside LLM-as-a-judge scoring for instruction-following and reasoning tasks.
- **Kimi-Eval** (Ding et al., 2025): A multilingual and multi-model evaluation suite designed to assess leading Chinese and English large language models, including the Baichuan series, Qwen, GLM, and Kimi itself. The benchmark spans automatic speech recognition (ASR), multiple choice question answering (MQA), open question answering (OpenQA), and reference-based question answering (RefQA), enabling a broad assessment of both comprehension and generative audio capabilities.
- **VoiceBench** (Chen et al., 2024): A focused benchmark evaluating thirty-five-plus state-of-the-art speech models across seven carefully selected datasets. While the total number of datasets is smaller than in AudioBench, the high task complexity and distinctive challenge of each dataset provide a useful test suite.
- **LMMS-Eval** (Zhang et al., 2025b): A comprehensive evaluation kit designed to assess multimodal frontier models across vision, audio and video modalities. Despite its broad coverage, audio-centric evaluation is comparatively limited as compared to other modalities.
- **AHELM** (Lee et al., 2025): A holistic evaluation benchmark aggregating diverse datasets across ten evaluation aspects to comprehensively assess audio-language models. Beyond presenting itself as benchmark, AHELM also functions as an evaluation kit, offering limited support for custom inference configurations and scalable throughput.

	MELD-Emotion	Librispeech-clean	ClothoAQA
# Samples	500	500	500
Audio Duration (seconds)	1,476	3,780	11,376
Evaluation Metric	LLM-Judge	WER	LLM-Judge

Table 6: **Experimental setup for efficiency comparison across evaluation frameworks.** We conduct controlled experiments using 500 samples from three diverse datasets: MELD-Emotion (short emotional speech), LibriSpeech-clean (medium-length read speech), and ClothoAQA (long-form descriptive audio). Total audio duration varies from 1,476 to 11,376 seconds, enabling assessment across different audio characteristics and evaluation modalities (LLM-judge vs. traditional metrics).

A.4 RTF and SPS Formulation Details

As discussed in Section 3, SPS and RTF are essential metrics to measure the efficiency of the evaluation framework. While RTF measures the processing time of an evaluation framework relative to the duration of the processed audio (Arriaga et al., 2024), SPS directly quantifies the model’s processing speed by measuring the average number of audio samples processed per second. Lower RTF is more desirable, and higher SPS is more preferable. Both metrics are formalized as follows:

$$RTF = \frac{\sum_{i=1}^N T_{i,proc}}{\sum_{i=1}^N D_{i,audio}}, \quad SPS = \frac{N}{\sum_{i=1}^N T_{i,proc}} \quad (1)$$

where T_{proc} is the total time (in seconds) taken by the framework to process the evaluation of the given audio. D_{audio} is the total duration of the input audio signal (in seconds) under 16kHz sampling rate, N is the total number of audio samples processed.

A.5 Inference Efficiency Ablations

To assess the scalability and efficiency of AU-Harness, we conduct three controlled ablations: (a) varying batch size, (b) throughput gains from parallel execution, and (c) latency trade-offs with replica scaling. The experimental setup follows Table 6, except for (c), where we use the full LibriSpeech-clean dataset to ensure sufficient workload for scalability analysis.

Figure 6 presents the results. Increasing batch size reduces execution time substantially, though benefits taper off at higher scales. Parallel execution yields up to a $3.5\times$ improvement in throughput over sequential execution, confirming the efficiency of concurrent scheduling. Replica scaling further lowers latency, with near-linear improvements observed up to 25 replicas.

Overall, these ablations highlight that AU-Harness is both scalable and adaptable. By leveraging batching, parallelism, and replica scaling, it can be tuned for diverse deployment scenarios ranging from high-throughput evaluation to low-latency inference.

A.6 Instruction Following Experimental Settings

This section provides complete transparency on the configurations and evaluation protocols used for each evaluation paradigm presented in Table 4. All conditions operate on the same audio-format inputs sourced directly from the original benchmark releases — no text-to-speech synthesis was performed at any stage. To ensure fair and reproducible comparison, all conditions share identical decoding parameters: `max_new_tokens=4096`, `temperature = 0.0001` (near-greedy decoding), and no system prompt unless otherwise specified. The maximum output token length is intentionally set high to accommodate the internal thinking traces emitted by Qwen3-Omni in thinking mode; these traces are systematically removed via post-processing before any performance metric is computed, ensuring that only the final response is evaluated. Without the loss of generality, we centralize our detailed explanations of the settings on the exemplar reasoning-focused task of BigBenchAudio. Additionally, we present our analysis on specific Qwen3-Omni (Xu et al., 2025b) model as there exists a dedicated ASR-specialized variant (Qwen3-Captioner). This enables a uniquely controlled three-tier transcription quality comparison: (1) model’s own internal ASR (Qwen3-Omni), (2) a same-family specialized transcription model (Qwen3-Captioner), and (3) the state-of-the-art external system (Whisper-large-v3 (Radford et al., 2023)). This design isolates the potential effects of transcription quality on downstream task performance without introducing confounding cross-family architectural differences. This study enables researchers to systematically explore and analyze the speech-text modality gap phenomenon and potential impacts of cross-modality interactions over downstream tasks as documented in prior works (Xiang et al., 2025; Cuervo et al., 2026; Nguyen et al., 2023, 2024, 2025). The evaluation paradigm is nonetheless model-agnostic and can be readily applied to any LALMs.




Figure 6: **Inference efficiency ablations in AU-Harness.** We examine three factors: (a) impact of batch size on execution time, (b) throughput gains from parallel execution, and (c) latency reduction through replica scaling.

Text Baseline

The text baseline feeds the original benchmark text directly to the model without any audio input, using the standard task prompt provided by each benchmark. This serves as the upper-bound reference for all audio-based conditions.

Answer the following reasoning question with a single word or number only – for example: "Yes", "No", "valid", "invalid", or a number. Do not explain your reasoning.

Begin by stating the question, then provide your single-word or single-number answer.

Prompt-only: Direct Audio

The model receives the raw audio file as input alongside the standard benchmark task prompt, with no additional instructions regarding transcription or intermediate steps. The model is expected to process the spoken content and produce a response end-to-end in a single inference pass. The full prompt template is as follows:

You are given an audio recording containing a reasoning question. Listen carefully and respond with a single word or number only – for example: "Yes", "No", "valid", "invalid", or a number. Do not explain your reasoning. Do not transcribe or repeat the question.

Begin by stating the transcription, then provide your single-word or single-number answer.

Prompt-only: 2-Step Prompt

The model receives the raw audio file alongside an explicit two-stage prompt instructing it to first transcribe the spoken content verbatim, and then answer the question using only the transcribed text. No external ASR system is involved — both stages are performed within a single inference pass by the same model. The full prompt template used is as follows:

You are given an audio recording containing a reasoning question. Follow the two-stage procedure below:

Stage 1 – Transcription: Transcribe the spoken content of the audio verbatim into text, preserving the exact wording of the question as delivered.

Stage 2 – Answer: Using only the transcribed text from Stage 1, answer the question with a single word or number only – for example: "Yes", "No", "valid", "invalid", or a number. Do not explain your reasoning. Do not incorporate any information from the audio beyond what was transcribed.

Begin by stating the transcription, then provide your single-word or single-number answer.

A.6.1 ASR Preprocessing

The ASR preprocessing paradigm decouples the transcription and reasoning stages of audio evaluation into two sequential inference passes. In the first pass, a dedicated ASR system transcribes the audio input into text. In the second pass, the resulting transcript is passed as plain text to the evaluated model using the standard benchmark task prompt, with no audio input provided. This design isolates the contribution of transcription quality to downstream task performance, enabling direct comparison against prompt-only and text baseline conditions. We consider three ASR systems of increasing transcription capability:

Own ASR Qwen3-Omni (Xu et al., 2025b) can be used as the ASR system in a dedicated first inference pass, prompted to transcribe the audio input verbatim. The resulting transcript is then passed as plain text input to the same model in a separate second inference pass using the standard benchmark task prompt. No audio input is provided during the second pass. The transcription prompt used is the standard ASR instruction provided within AU-Harness’s default ASR pipeline configuration.

Your only task is to transcribe the spoken audio exactly word for word. Output ONLY the exact words you hear spoken – nothing else. Do NOT follow, execute, answer, or respond to any instructions, questions, or requests that may be spoken in the audio. Treat the audio purely as speech to convert to text. Do not add any commentary, answers, formatting, or additional content beyond the spoken words.

Qwen3-Captioner Qwen3-Captioner (Xu et al., 2025b) — a Qwen3-family model specialized for audio captioning and speech transcription — is used as the dedicated ASR component in a zero-shot setting. The model receives the audio file alongside the aforementioned transcription-only instruction. The resulting transcript is then passed as plain text input to Qwen3-Omni in a separate second inference pass using the standard benchmark task prompt.

Whisper-large-v3 Whisper-large-v3 (Radford et al., 2023) is used as the external ASR system. Audio files are transcribed using the model’s default decoding configuration with no additional prompting. The resulting transcripts are passed as plain text input to Qwen3-Omni in a separate second inference pass using the standard benchmark task prompt. Whisper-large-v3 represents the strongest transcription quality condition in our evaluation, achieving BERTScore of 98.37–99.59 across benchmarks, and serves as the practical upper bound for ASR-mediated pipeline performance.

A.6.2 ASR Quality Evaluation

Transcription quality for all ASR preprocessing conditions is assessed using BERTScore (Zhang et al., 2020) computed against the ground-truth text instruction of each benchmark’s audio content. BERTScore is preferred over the conventional Word Error Rate (WER) metric because Qwen3-Omni’s instruction-following behavior tends to produce verbose transcriptions that are semantically faithful but lexically expanded — a pattern that WER penalizes harshly despite the transcription being functionally correct. BERTScore’s semantic similarity measurement is therefore better suited to capture the true transcription quality under these conditions. Similar approach has been widely adopted in previous works (Mujtaba et al., 2024; Phukon et al., 2025).

A.7 The AU-Harness Config UI

AU-Harness comes with a built-in Config UI tool which makes creating the config for setting up runs

easier (Figure 7). The output is a generated config YAML which can be checked by a researcher before submitting the job.

A.8 The AU-Harness Run Report

The AU-Harness run report is a self-contained HTML artifact constructed at the end of every evaluation run. It surfaces the same underlying records used to populate the leaderboard, but reorganizes them around the questions practitioners typically ask once a run completes: which categories went well, where did regressions land, what failed at the infrastructure layer, and long-context behavior analysis for multi-turn benchmarks. The report exposes seven tabs, summarized below; Figures 8, 9, 10, 11, 12, 13, 14 show each one for a run config with Qwen2.5-Omni-3B and 7B, Voxtral-Mini-3b, and Voxtral-Small-24B.

A.9 Turn Range Selection

We cap evaluation at 20 turns because the all-audio configurations (where each user turn is raw audio alongside agent turn as text) approach the context length limits of current audio language models, and can lead to truncation beyond this point. While text-only configurations could in principle extend further, we fix the window across all settings to ensure comparable and consistent evaluation. Additionally, we omit the first turn from all plots as the opening turn is nearly always a greeting that introduces no new slots and adds visual noise without analytical value, especially for SpokenWOZ.




A.10 Licensing

AU-Harness is released under Apache-2.0 license. The framework does not host or contain datasets; therefore, running evaluations of certain datasets with AU-Harness follow the licenses of the original datasets’ publishers and the aforementioned AU-Harness’ license.

A.11 User Case Study

To assess the ease-of-adoption of our AU-Harness, we conduct a controlled user case study. More specifically, our study involves 5 participants experienced in LLM evaluations and the users are expected to leverage AU-Harness and extend it to new tasks and model from scratch. Following Joshi et al. (2025); Ismanto et al. (2026), we compute System Usability Scale (SUS) and Technology Acceptance Model (TAM) assessment scores.

Participant Background All selected participants had prior hands-on experience with at least one of the contemporary LLM evaluation toolkits.






Figure 10: **Categories.** Categories groups every (dataset, metric, model) result under one of the AU-Harness category labels — Audio Understanding, Paralinguistics, Safety and Security, Speech Disorder / Enhancement, Speech Recognition, Spoken Language Reasoning, and Spoken Language Understanding. Within each block, rows are sorted by score and color-coded against the category mean; each block closes with a category-average row aggregating over the (dataset × model) entries in scope. The view is intended to expose category-level strengths that a flat leaderboard obscures.




Figure 11: **Patterns.** The Patterns tab is the per-task error-mining surface. For each (task, model) pair, AU-Harness partitions samples into low-score and high-score buckets (counts shown in the header) and computes the prompt terms and judge-rationale phrases via LDA that are disproportionately overrepresented in the low bucket. Low-score prompt terms are tokens drawn from the dataset’s user-side instruction; judge low-score phrases are tokens from the judge model’s free-text justification on failing samples. Together they provide a coarse but fast signal of where a model is breaking down — vocabulary it cannot ground, content it refuses, output format it mis-aligns - without requiring per-sample inspection.




Figure 12: **Multi-turn.** Multi-turn targets long-conversation degradation, currently instrumented on SpokenWOZ-style dialogue state tracking with joint goal accuracy, slot accuracy, and slot F1. For each (task, metric, model) triple the tab plots mean score per turn, summarizes early (turns 0–3), middle (4–9), and late (10+) buckets, and renders a failure-turn histogram showing which turn indices contribute the most zero-score samples. Bucket cells are tinted by absolute level, making floor effects in the late bucket visually distinguishable from a smooth glide.

This decision is deliberate to ensure the fair comparative assessment of the AU-Harness’ usability.

Guideline for participants We evaluate AU-Harness through the following realistic workflows when utilizing AU-Harness and extending it to new task(s) and new model(s).

1. Install and configure the toolkit directly from source code
2. Set up the evaluation environment by reviewing README documentation (specifically *run_config*, *task_config* definitions and specifications)
3. Create *task_config* to integrate the specified task into the toolkit
4. Introduce the model configurations to the *run_config*
5. Execute the benchmark with the toolkit
6. Inspect result logs and validate successful completion

To better the scope the study, the target model and task are specified upfront, enabling the focus on integrating them into AU-Harness. For task, we select MMAU-PRO (music subset) (Kumar et al., 2026) and for the model, we leverage Qwen3-Omni-30B-A3B-instruct (Xu et al., 2025b).

Observations The observed results are strongly positive among participants of AU-Harness user study. SUS score is averaged 87.5/100 (an excellent rating above the industry benchmark) indicating participants found AU-Harness genuinely easy

AU-Harness Run Report — voxtral3b_voxtral24b_gwen2p5-3b
Model 1: infer-open-2p5-omni-3b, Model 2: infer-open-2p5-omni-7b, Model 3: infer-voxtral-mini-3b, Model 4: infer-voxtral-small-24b

Dataset	Metric	Dir	infer-open-2p5-omni-3b	infer-open-2p5-omni-7b	infer-voxtral-mini-3b	infer-voxtral-small-24b
noise_detection	F1_L1_arity	↑	59.00	—	12.50	15.50
audiobench_test	F1_L1_arity	↑	48.76	—	26.84	32.20
voicebench_test_audio	Instruction_L1_bearing	↑	32.46	—	—	59.71
ad_bench	F1_L1_arity	↑	68.70	—	88.70	86.40
big_bench_hard_audio_query	F1_L1_arity	↑	33.20	—	46.60	58.00
siemens_detection	F1_L1_arity	↑	53.80	—	2.00	49.80
iemocap_gender_recognition	F1_L1_arity	↑	69.90	—	46.70	62.60
mmmi_test	F1_L1_arity	↑	53.30	—	40.60	48.10
briseptest_test_mean	word_error_rate	↓	13.01	—	2.42	2.03
mlu_promis_test	F1_L1_arity	↑	51.70	—	45.30	58.60
spokenword_audio_test	join_L1_accuracy	↑	—	9.20	10.70	—
spokenword_audio_test	MLU_accuracy	↑	—	36.49	35.53	—
spokenword_audio_test	MLU1	↑	—	37.54	37.54	—
spokenword_audio_response	llm	↑	—	2.45	1.35	—
spokenword_audio_test	join_L1_accuracy	↑	—	11.10	12.39	—
spokenword_audio_test	MLU_accuracy	↑	—	41.54	40.94	—
spokenword_audio_test	MLU1	↑	—	43.34	43.99	—
spokenword_audio_response	llm	↑	—	2.61	1.35	—
spokenword_audio_test_sum	join_L1_accuracy	↑	—	32.20	37.59	—

Figure 13: **Compare.** Compare is a head-to-head view that accepts up to four model selections and renders per-task scores side by side. A Dir column encodes metric direction (↑ higher is better; ↓ lower is better, e.g. word error rate), and the leading score per row is bolded. The header line summarizes wins by aggregating direction-aware comparisons across the tasks where every selected model has a valid score; tasks with any missing score are excluded from the win count but remain visible in the table.

to install, configure, and operate despite the complexity of the workflow. TAM score is averaged 4.67/5, reflecting strong perceived usefulness, ease of use, and a clear intention to adopt and recommend the tool for future LALM evaluations. Together, these scores suggest AU-Harness delivers a smooth setup experience and real research value, with participants encountering few usability barriers when extending it to new models and benchmarks. Because the study tests the tool rather than the participants, any friction reported is treated as actionable feedback for improving the tool design.

AU-Harness Run Report — voxtral3b_voxtral24b_gwen2p5-3b
Model 1: infer-open-2p5-omni-3b, Model 2: infer-open-2p5-omni-7b, Model 3: infer-voxtral-mini-3b, Model 4: infer-voxtral-small-24b

Dataset	Metric	Model	Total	Failures	Failure rate	Avg wait	Top error
mmmi_test	F1_L1_arity	infer-open-2p5-omni-3b	1000	91	9.1%	15.70s	error_server
voicebench_test_audio	Instruction_L1_bearing	infer-open-2p5-omni-3b	345	28	8.1%	28.37s	error_server
ad_bench	F1_L1_arity	infer-open-2p5-omni-3b	520	39	7.5%	11.79s	error_server
audiobench_test	F1_L1_arity	infer-open-2p5-omni-3b	1000	66	6.6%	13.33s	error_server
mlu_promis_test	F1_L1_arity	infer-open-2p5-omni-3b	1187	67	5.6%	13.54s	error_server
big_bench_hard_audio_query	F1_L1_arity	infer-open-2p5-omni-3b	1000	39	3.9%	13.51s	error_server
briseptest_test_mean	word_error_rate	infer-open-2p5-omni-3b	2617	95	3.6%	17.22s	error_server
iemocap_gender_recognition	F1_L1_arity	infer-open-2p5-omni-3b	1004	31	3.1%	13.88s	error_server
noise_detection	F1_L1_arity	infer-open-2p5-omni-3b	200	0	0.0%	10.77s	n/a
audiobench_test	F1_L1_arity	infer-open-2p5-omni-3b	313	0	0.0%	10.67s	n/a
voicebench_test_audio	Instruction_L1_bearing	infer-voxtral-mini-3b	345	0	0.0%	6.79s	n/a
ad_bench	F1_L1_arity	infer-voxtral-small-24b	520	0	0.0%	4.52s	n/a
audiobench_test	F1_L1_arity	infer-voxtral-mini-2b	520	0	0.0%	2.15s	n/a
audiobench_test	F1_L1_arity	infer-voxtral-small-24b	313	0	0.0%	2.08s	n/a
mlu_promis_test	F1_L1_arity	infer-voxtral-small-24b	1187	0	0.0%	2.07s	n/a
mmmi_test	F1_L1_arity	infer-voxtral-small-24b	1000	0	0.0%	1.47s	n/a
briseptest_test_mean	word_error_rate	infer-voxtral-small-24b	2617	0	0.0%	1.43s	n/a
iemocap_gender_recognition	F1_L1_arity	infer-voxtral-small-24b	1004	0	0.0%	1.32s	n/a
audiobench_test	F1_L1_arity	infer-voxtral-mini-3b	313	0	0.0%	1.23s	n/a
noise_detection	F1_L1_arity	infer-voxtral-small-24b	200	0	0.0%	1.10s	n/a
spokenword_audio_response_summarized	llm	infer-voxtral-mini-3b	1000	0	0.0%	1.05s	n/a
spokenword_audio_test_summarized	join_L1_accuracy	infer-voxtral-mini-2b	1000	0	0.0%	0.96s	n/a

Figure 14: **Ops Health.** Ops health separates inference reliability from model quality. For each (dataset, metric, model) row it reports the total request count, hard-failure count, failure rate, mean wall-clock wait per request, and the most common error class. Highlighted rows surface transient API outages, rate-limit storms, or per-model timeout misconfigurations that would otherwise be mis-read as quality regressions in the other tabs.